LLNL-TR-408454

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Preliminary Neutronics Design Studies for a Molten Salt Blanket LIFE Engine

J. Powers

November 6, 2008

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

August 6, 2008

Preliminary Neutronics Design Studies for a Molten Salt Blanket LIFE Engine

Jeff Powers
Engineering and NIF Summer Scholar
Department of Nuclear Engineering
University of California – Berkeley

TABLE OF CONTENTS

1. INTRODUCTION AND BACKGROUND

The Laser Inertial Confinement Fusion Fission Energy (LIFE) Program being developed at Lawrence Livermore National Laboratory (LLNL) aims to design a hybrid fission-fusion subcritical nuclear engine that uses a laser-driven Inertial Confinement Fusion (ICF) system to drive a subcritical fission blanket. This combined fusion-fission hybrid system could be used for generating electricity, material transmutation or incineration, or other applications. LIFE does not require enriched fuel since it is a sub-critical system and LIFE can sustain power operation beyond the burnup levels at which typical fission reactors need to be refueled. In light of these factors, numerous options have been suggested and are being investigated. Options being investigated include fueling LIFE engines with spent nuclear fuel to aid in disposal/incineration of commercial spent nuclear fuel or using depleted uranium or thorium fueled options to enhance proliferation resistance and utilize non-fissile materials [1].

LIFE engine blanket designs using a molten salt fuel system represent one area of investigation. Possible applications of a LIFE engine with a molten salt blanket include uses as a spent nuclear fuel burner, fissile fuel breeding platform, and providing a backup alternative to other LIFE engine blanket designs using TRISO fuel particles in case the TRISO particles are found to be unable to withstand the irradiation they will be subjected to. These molten salts consist of a mixture of LiF with UF_4 or ThF_4 or some combination thereof. Future systems could look at using PuF_3 or PuF_4 as well, though no work on such system with initial plutonium loadings has been performed for studies documented in this report.

Figure 1 provides an illustration of what a complete LIFE engine might look like if driven by a fast ignition fusion system.

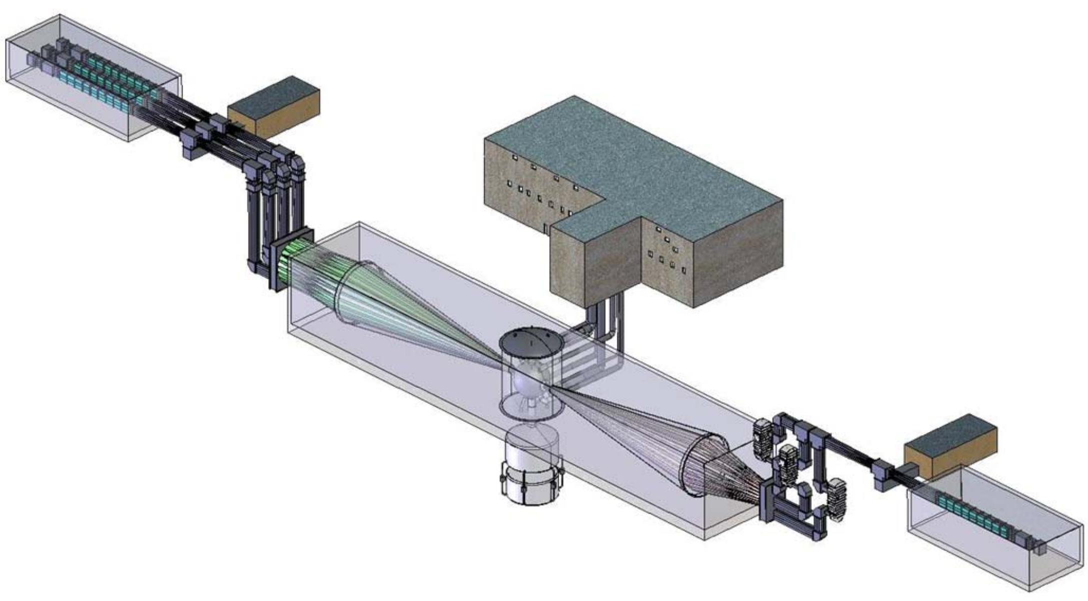

Figure 1. Illustration of a Fast Ignition LIFE Engine

2. PURPOSE AND SCOPE

The purpose of this report is to document preliminary neutronics design studies performed to support the development of a molten salt blanket LIFE engine option, as part of the LIFE Program being performed at Lawrence Livermore National laboratory. Preliminary design studies looking at fast ignition and hot spot ignition fusion options are documented, along with limited scoping studies performed to investigate other options of interest that surfaced during the main design effort. Lastly, side studies that were not part of the main design effort but may alter future work performed on LIFE engine designs are shown. The majority of all work reported in this document was performed during the Molten Salt Fast Ignition Moderator Study (MSFIMS) which sought to optimize the amount of moderator mixed into the molten salt region in order to produce the most compelling design.

The studies in this report are of a limited scope and are intended to provide a preliminary neutronics analysis of the design concepts described herein to help guide decision processes and explore various options that a LIFE engine with a molten salt blanket might enable. None of the designs shown in this report, even reference cases selected for detailed description and analysis, have been fully optimized. The analyses were performed primarily as a neutronics study, though some consultation was made regarding thermal-hydraulic and structural concerns during both scoping out an initial model and subsequent to identifying a neutronics-based reference case to ensure that the design work contained no glaring mechanical or thermal issues that would preclude its feasibility. Any analyses and recommendations made in this report are either primarily or solely from the point of view of LIFE neutronics and ignore other fundamental issues related to molten salt fuel blankets such as chemical processing feasibility and political feasibility of a molten salt system.

3. METHODOLOGY

As with other neutronics calculations being performed in support of the LIFE Program, the studies documented in this report used the MONTEBURNS 2.0 neutronics burnup code package [2], which couples MCNP5 [3] for neutron and radiation transport to ORIGEN-2.2 for isotopic evolution due to burnup and decay. The reasons for using Monte Carlo based neutron and radiation transport, and more specifically the MONTEBURNS 2.0 burnup code package, have been documented previously and are seen as prudent for this stage of the LIFE Program [4].

In addition to the MONTEBURNS 2.0 burnup code package, other codes and software were used during the course of the analyses in this report. Of particular importance was TBR [5], a code written by Kevin Kramer at LLNL with contributions from Jeff Latkowski, which interfaced with MONTEBURNS 2.0 to simplify and optimize necessary adjustments to material compositions being used in MONTEBURNS calculations. These adjustments were necessary to balance user-specified design parameters such as maximum power level for the LIFE engine to run at, acceptable tritium breeding ratio (the parameter for which the TBR code is named) ranges during operation, and a host of options which enable constant removal or addition of various isotopes between MONTEBURNS calculations. The TBR code allowed for an iteration scheme to ensure that the LIFE engine modeled by MONTEBURNS followed the specifications and constraints that the user had in mind for that particular concept, as well as some post-processing of results to condense the many output files from MONTEBURNS into single concise files. In addition to the above discussion about the TBR code, it should be noted that all MCNP calculations performed during the analyses in this report used the ".70c" MCNP cross section data library generated by Peter Song for the LIFE Program. All MCNP calculations were performed using cross sections generated for a temperature of 600°C and moderating materials were treating using free gas treatment rather than bound-atom S(α,β) thermal scattering treatment in MCNP. These approximations to cross section temperatures and thermal neutron scattering treatment are considered acceptable for a preliminary neutronics design study.

Further analysis of results was performed manually by the analyst using shell scripts in Linux, macros in the Vim text editor in Linux, and Microsoft Excel for data manipulation and visualization. Microsoft Excel was also used for MCNP geometry and material input calculations.

Figure 2 shows a sketch of a typical curve with the LIFE engine power output in megawatts of thermal power (MW_{th}) as a function of time during a complete operational lifetime for a LIFE engine with an initial actinide loading of fertile materials such as depleted uranium and thorium.

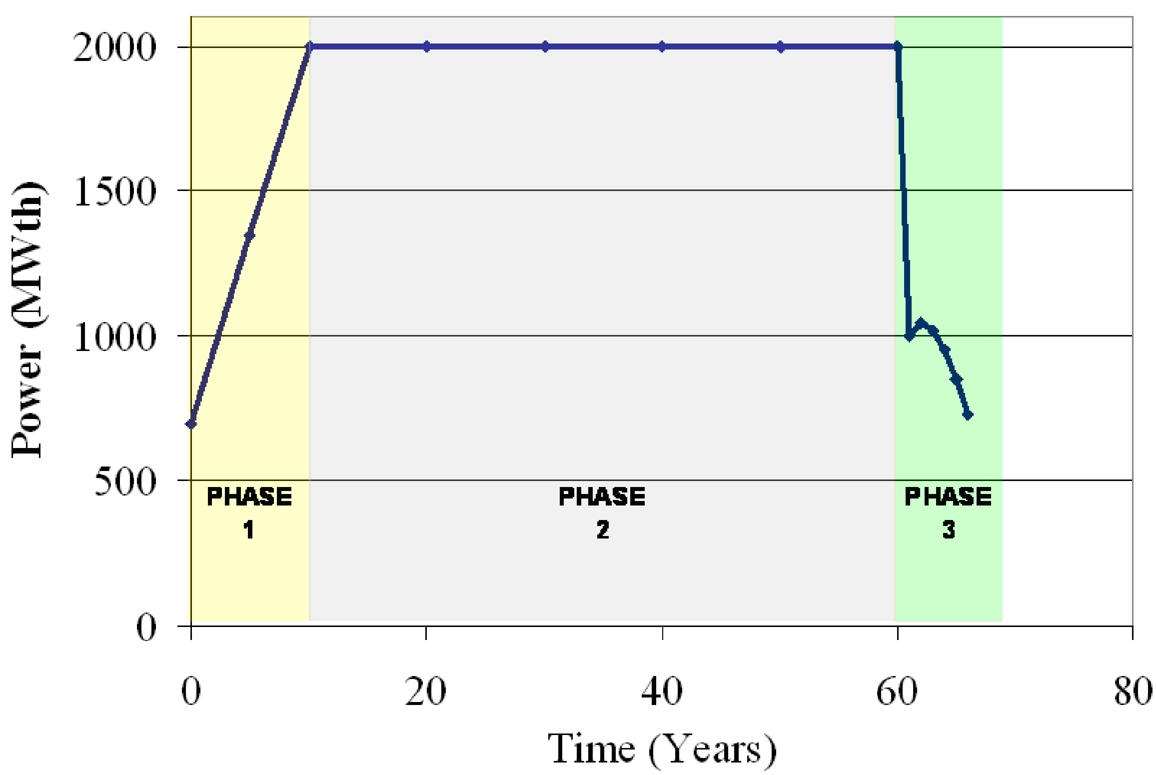

Figure 2. Sample Plot of LIFE Engine Power Output During Operation

The power curve has been divided into three distinct phases which can be understood separately. During Phase 1, power output climbs as neutrons in the system are used to fission existing material while also converting fertile materials (e.g., Th-232 and U-238) into fissile materials (e.g., U-233 and Pu-239). Phase 2 begins when these fertile materials have reached sufficient levels to sustain the thermal power output necessary for the user-specified plateau power level. In this case, Phase 2 begins 10 years into operation of the LIFE engine at a plateau power level of 2000 MW_{th}. The plateau power level is maintained during Phase 2 by using Li-6 to absorb extra neutrons, thus decreasing the number of fission events while also breeding tritium for use in the fusion driver. The use of lithium in this manner is similar to the use of burnable poisons or other reactivity control mechanisms in nuclear fission reactors but has the added benefit that tritium is produced and can be used as fusion fuel at some later point during operation or for another application (e.g., a different LIFE engine). Phase 3, which begins 60 years into operational life in Figure 2, indicates a shift out of power operation mode. This shift may be due to exhausting the existing tritium inventory, in which case the LIFE engine switches into tritium breeding mode and varies the thermal power output as needed to maintain a user-specified tritium breeding ratio (TBR). A shift out of Phase 2 may also be due to an insufficient inventory of fissile and fissionable material to maintain the fission rate needed for the plateau power level. The curve ends during Phase 3 when it reaches the specified number of burnup time steps or the analyst can cut the curve at a desired point in operation, such as a burnup level target.

During the course of these analyses, it was determined that the definition of performance indices would be useful in providing some standardization and equality in evaluating concepts that might be of very different design. To this end, two performance indices were defined and are reported for various concepts in this report: the "Plateau MWD Fraction" index and the "Plateau Time Fraction" performance index, both of which are described below and would have values of 1.0 for an idealized LIFE engine design. Figure 3, below, shows sketches of both idealized and more realistic power curves for LIFE engine designs and will be referred to in the following paragraphs to help illustrate the performance index definitions. It is important to note that since all burnup calculations in this report were performed with 90 day time steps, performance indices do not have perfect resolution and thus should not be used to decide between two cases with very small differences in indices. The use of 90 day time steps has had no explicit verification and validation but has been used in previous LIFE Program studies.

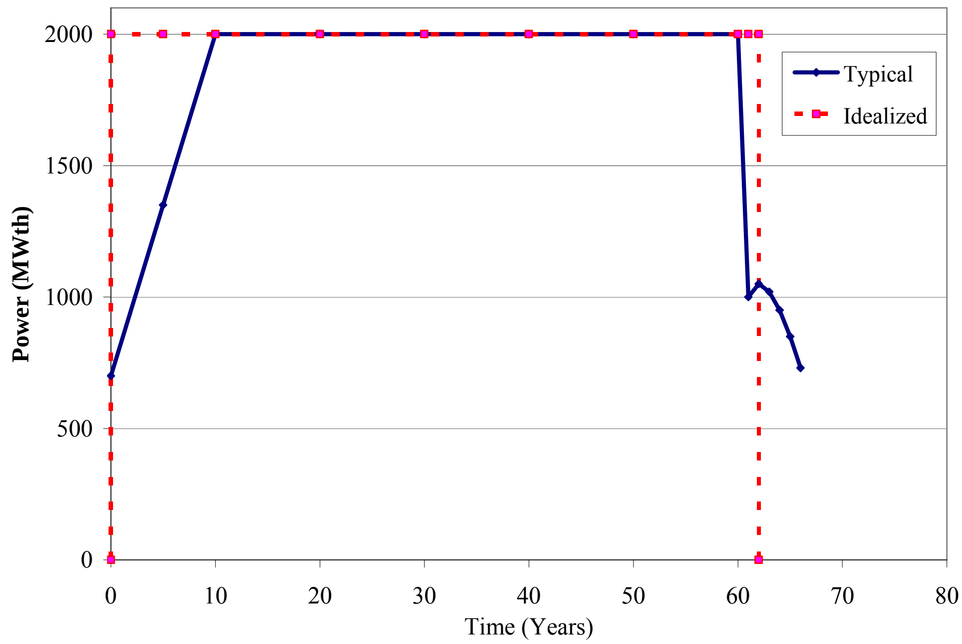

Figure 3. Power Curves for Fictional "Typical" and "Idealized" Cases

The "Plateau MWD Fraction" performance index provides a measure of what fraction of total energy (in units of megawatt days or MWD for short) produced during a LIFE engine's lifetime occurs while at the peak design power. This can be explained visually using Figure 3 and noting that the Plateau MWD Fraction would be calculated as the integral of the entire power curves divided by the integral of the portion of the power curve at the plateau (i.e., from 10 years to 60 years for the "Typical" curve). Since the "Idealized" case exists entirely as a plateau it would have an index of 1.0, while the "Typical" Curve here would have an index of about 0.9.

The "Plateau Time Fraction" performance index provides a measure of what fraction of total operational time of a LIFE engine, from beginning of life (BOL) to end of life (EOL), occurs while at the peak design power. EOL must be defined by the analyst for a given study or concept and may be defined different depending on the mission of the LIFE engine being analyzed; as will be discussed in further detail later, the EOL criterion used for all analyses in this report was 99% burnup of all actinides. Mathematically, this is calculated as the time difference between the time step right after the plateau and the first time step of the plateau divided by the time of the first time step with over 99% FIMA. For Figure 3, the "Idealized" case would have a Plateau Time Fraction index of 1.0 while the "Typical" case would have a Plateau Time Fraction index of only about 0.77.

4. CONCEPT DESCRIPTION

A preliminary study was performed by Peter Song at LLNL to investigate various options in geometry and molten salt composition for a molten salt fast ignition LIFE design. These molten salt studies were performed without any moderator in the molten salt region. Different molten salt compositions with different levels of UF_4 and ThF_4 were investigated in order to understand the tradeoffs between breeding fissile materials in the form of U-233 and Pu-239 while limited the peak plutonium concentration in the molten salt to acceptable levels for solubility and proliferation concerns. Peter Song identified case F55, which had a molten salt of 12 mol% UF_4 and 12mol% ThF_4 with 76 mol% LiF, as the most promising option of those he studied. Case F55 was therefore taken as the initial case for the molten salt fast ignition moderator study (MSFIMS).

All molten salt studies documented in this report assumed the continuous active removal of various noble and semi-noble metals and gaseous fission products. The elements being removed were: Kr, Xe, Zn, Ga, Ge, As, Nb, Mo, Tc, Ru, Rh, Pd, Ag, Cd, In, Sn, and Sb. In addition, as with all burnup calculations performed for the LIFE Program, tritium was continuously removed from the system as well to provide new fusion fuel and minimize potential biological or engineering problems due to tritium contamination.

In addition, the molten salt studies assumed that the number of fluorine atoms was constantly being reset to the initial number of fluorine atoms present at the beginning of operation. This is not accurate since fluorine atoms may be produced or destroyed during operation. More importantly, the number of fluorine atoms in the molten salt may need to be modified to mitigate corrosion effects or handle the chemistry effects that stem from initial uranium and thorium atoms turning into multiple fission products. These issues should be examined and quantified in future studies.

An illustration of a complete fast ignition LIFE engine can be found in Figure 1. A picture of the actual MCNP model, obtained using VISED [5], can be found along with some basic component labels in Figure 4. The unlabelled regions surrounding the beryllium multiplier region in the picture are the coolant plena described later in Table 1. Note that, as pictured, each region is fully homogenized; this approximation seems reasonable for any preliminary neutronics studies, but should prove especially reasonable for a molten salt case where the fuel truly is nearly homogenous.

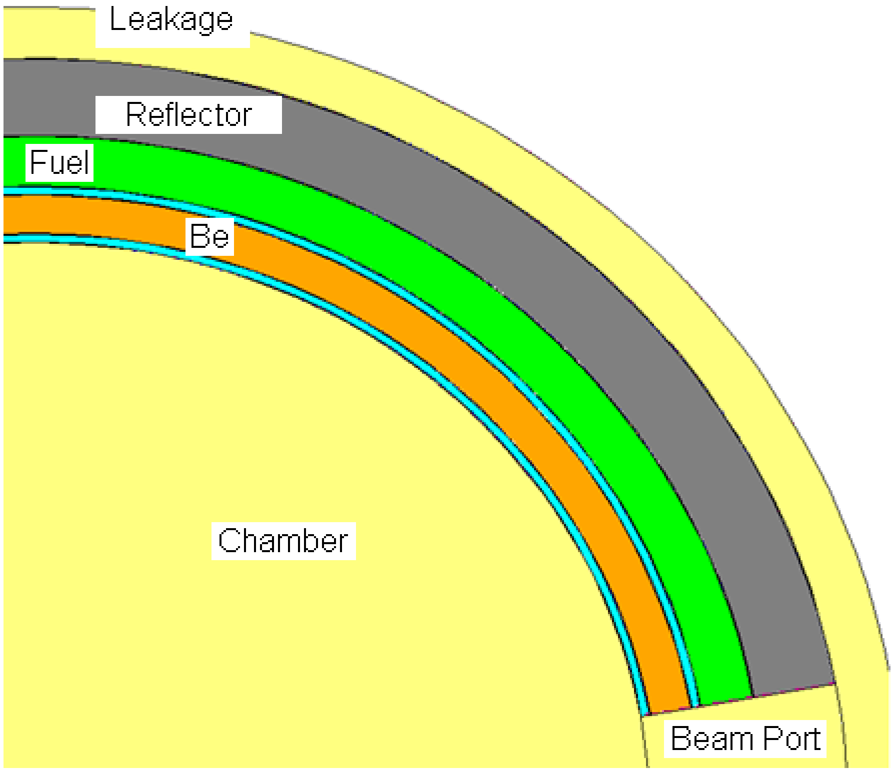

Figure 4. VISED-Produced Picture of MCNP Model for Case F55

5. DESIGN OBJECTIVES AND CONSTRAINTS

The primary objectives for the analyses in this report were reaching 99% FIMA burnup, minimizing the incineration time, and maximizing energy production at the peak plateau power level.

Burnup is often measured in units of Fraction of Initial Metal Atoms (FIMA) consumed or %FIMA. This parameter represents the fraction, or percent, of initial actinide atoms consumed during operation. Incineration time, defined as the time required after the start of operation for a design to reach the 99% FIMA burnup goal, provides a measure of how effective a design is at burning through the initial actinide mass. Minimizing this incineration time while maximizing the time spent at, and energy produced during, peak plateau power operation formed some of the fundamental design goals during the analyses in this report. Unless otherwise noted, all results shown in this report have their curves and data extended to 99% FIMA burnup and design decisions were made with the goal meeting the two objectives above.

The primary constraints assumed for all analyses documented in this report, unless otherwise noted are as follows:

- The peak plateau power level was set to 2000 MW_{th} [a range of 2000 MW_{th} to 2050 MW_{th} in the “tbr.inp” input file]
- The initial fuel loading was 40 metric tons of actinides.
- The tritium breeding ratio (TBR) must be maintained between 0.3 and 3.0

- All dimensions, materials, and densities must remain consistent with the base case for that study other than the variables established for the study and any other parameters directly dependent on those variables (e.g., reflector radii move with the fuel region).
- A 30cm thick graphite reflector was used.

6. MODEL DESCRIPTION

This section describes the MCNP models used for the subsequent analyses in this report.

6.1. Molten Salt Fast Ignition Model

As mentioned previously, the initial base case used to begin work on the MSFIMS was Peter Song's case F55. The F55 model was therefore used as a starting for any model modifications and changes that were needed. Table 1 provides a listing of the material, density, and thickness of each major component of the F55 model.

Table 1. Material, Density, and Thickness Specifications for Components of Case F55

Component	Material	Density [g/cc]	Thickness [cm]
Chamber	Vacuum	1.29E-09	225 (IR)
First Wall	ODS Steel	8.0	0.3
Plenum	FLiBe	2.0	3
2nd Wall	Porous ODS Steel	6.0	0.3
Be + FLiBe	60 vol% Be 40 vol% FLiBe	1.91	15
3rd Wall	Porous ODS Steel	6.0	0.3
Plenum	FLiBe	2.0	3
4th Wall	ODS Steel	8.0	0.3
Fuel Zone	Molten Salt Graphite	2.683	48.98
5th Wall	ODS Steel	8.0	0.3
Reflector	Graphite	2.0	30
6th Wall	ODS Steel	8.0	0.3

The model described in Table 1 and pictured in Figure 4 needed modifications for the moderator study. Each new case in the moderator study needed several parameters changed: the fuel material density specified in the MCNP cell card, the fuel material composition specified in the MCNP data card, and the radii for all surfaces from the outer radius of the fuel region to the model boundary in the MCNP surface card. Graphite was added to the fuel region in a way that allowed the user to specify the desired volume percent of the total fuel and graphite region occupied by graphite while conserving the initial actinide loading and molten salt density and relative weight fractions. This expanded the fuel and graphite region outward as graphite was added and also pushed out the surrounding components while preserving their initial thicknesses. The molten salt and graphite were modeled as a single homogenous region for all of the studies in this report to simplify modeling efforts and avoid issues that arose in the MONTEBURNS and TBR codes when using an explicitly heterogeneous geometry model with multiple fuel regions separated by graphite shells.

6.2. Molten Salt Hot Spot Ignition Model

The initial model used as a basis for the molten salt hot spot ignition (HSI) model was case bb0, which was a TRISO-based HSI model. After speaking to Ryan Abbott at LLNL regarding the thermal and structural implications of using molten salt fuel instead of TRISO-based pebble fuel, a base case MCNP model for molten salt HSI was arrived at using 40tons of actinides. This model, identified as h10, is pictured below in Figure 5 and a full listing of the material, density, and thickness of each major component is listed in Table 2.

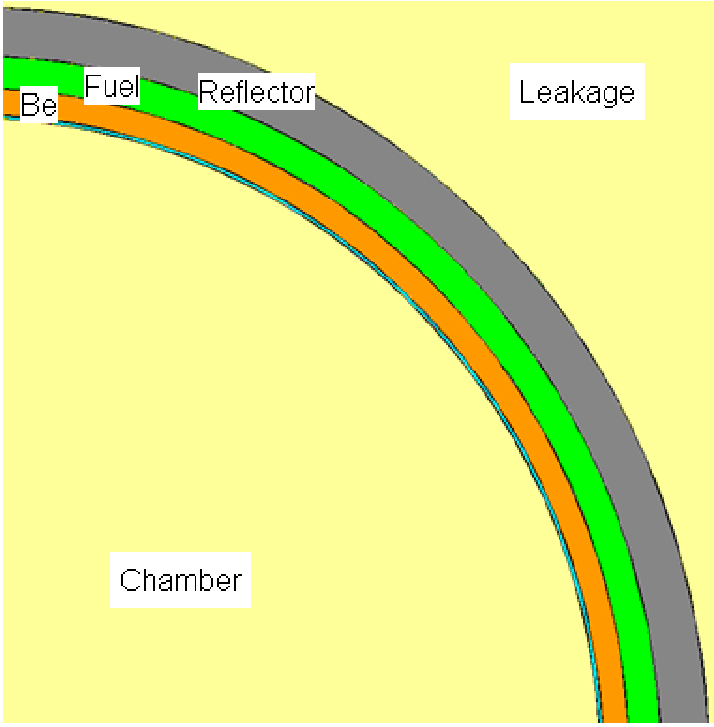

Figure 5. VISED-Produced Picture of MCNP Model for Case h10

Table 2. Material, Density, and Thickness Specifications for Components of Case h10

Component	Material	Density [g/cc]	Thickness [cm]
Chamber	Xenon	6.50E-06	400 (IR)
First Wall Armor	Tungsten	19.3	0.025
First Wall	ODS Steel	8.0	0.275
First Wall Coolant	FLiNaBe	2.0	2
2nd Wall	ODS Steel	8.0	0.3
Be + FLiBe	60 vol% Be 40 vol% FLiBe	1.91	16
3rd Wall	ODS Steel	8.0	0.3
Fuel Zone	Molten Salt Graphite	2.6831	19.715723
4th Wall	ODS Steel	8.0	0.3
Reflector	Graphite	2.0	30
5th Wall	ODS Steel	8.0	0.3

As with the fast ignition model, the molten salt and graphite were modeled as a single homogenous region for all of the studies in this report both to simplify modeling efforts and avoid issues that arose in the MONTEBURNS and TBR codes when using an explicitly heterogeneous geometry model that had multiple fuel regions separated by graphite shells. The actual molten salt composition used in case h10 70 vol% graphite and 30 vol% molten salt, with an identical composition to the molten salt fats ignition reference case (case G70).

The h10 base molten salt hot spot ignition model was modified as needed for other HSI cases by expanding the fuel and graphite region volumetrically and maintain the original thickness of regions outside of it. This expansion held constant the density of the molten salt and graphite region while also maintaining the composition as 70 vol% graphite and 30 vol% molten salt. Additional actinide mass was added by scaling the volume of the region to whatever volume necessary to yield the desired actinide mass.

7. MOLTEN SALT FAST IGNITION MODERATOR STUDY RESULTS

Results from the Molten Salt Fast Ignition Moderator Study (MSFIMS) performed during these analyses are shown below, along with a comparison of the fast ignition molten salt reference case concept (case G70) to the fast ignition TRISO fuel reference case concept (case da0). The reference case (G70) is named to represent the fact that it has a Graphite moderator which accounts for 70 vol% of the fuel and moderator region; this naming convention, cases being labeled as G## where ## is the vol% of moderator in that case, is used for all MSFIMS cases.

7.1. Molten Salt Results

The MSFIMS took the existing molten salt fast ignition base case, case F55, and performed a sweeping study with varying volume fractions of moderator in the molten salt region to identify the most promising range of moderator volume fraction to optimize the blanket design. Burnup calculations were performed with 0, 5, 10, 15, 20, 40, 50, 60, 70, 80, 90, and 95 vol% of moderator in the molten salt and moderator region. The series of figures below show the results of these burnup calculations; unless otherwise noted, all curves are stopped at 99% FIMA burnup to allow for fair comparisons.

It is important to note that all calculations have been found acceptable through some basic spot-checking and verification work as described later in Section 10, with minor exceptions for the G20 and G50 cases which appear to have small errors in the results. These cases appear to have been offset by one or two time steps or two due to errors in the burnup calculations, and thus their results differ slightly from what a fully accurate calculation would indicate. This offset is only a minor effect, however, and the overall trends identified during the MSFIMS hold true. Rerunning the calculations to obtain fully accurate results was considered but deemed unnecessary since the general trends are not affect and neither the G20 nor G50 case are in the region of most interest.

Table 3 provides the masses for each case in terms of actinide mass, total mass in the molten salt and graphite region, and total mass in the entire MCNP model of the system. It should be noted that this total system mass consists only of the components listed in Table 1, which does not include additional structure that will be needed or any of the equipment for the lasers or other fusion systems. Table 3 shows that the same initial actinide mass was used in each case, but as the moderator volume fraction increased the molten salt and graphite region increased nonlinearly in mass as did the total system mass. All masses are given in Metric Tons (MT).

Table 3. Important Masses for the MSFIMS

Case	Actinide Mass (MT)	Molten Salt + Graphite Mass (MT)	Total System Mass (MT)
F55	40.01	66.95	162.88
G05	40.01	68.15	164.45
G10	40.01	69.49	166.19
G15	40.01	70.99	168.14
G20	40.01	72.67	170.33
G40	40.01	82.19	182.69
G50	40.01	89.82	192.53
G60	40.01	101.25	207.23
G70	40.01	120.31	231.54
G80	40.01	158.42	279.63
G90	40.01	272.77	420.79
G95	40.01	501.45	694.95

The first results analyzed after burnup calculations have been completed are usually the power curves generated during the burnup calculations. The power curves provide a quick glance at how well the design performed, especially when compared to other similar designs to decide which performed better. Figure 6, shown below, provides the power curves for all calculations performed for the MSFIMS. The trends in power curves show the plateau length gradually increasing from 0 vol% moderator (case F55) to 70 vol% moderator (case G70), and then decreasing from 70 vol% moderator (case G70) to 95 vol% moderator (case G95). The power curve for G95 shows a clear decrease in performance compared to even F55 with no moderator. Based upon these curves, G70 offers the most compelling performance.

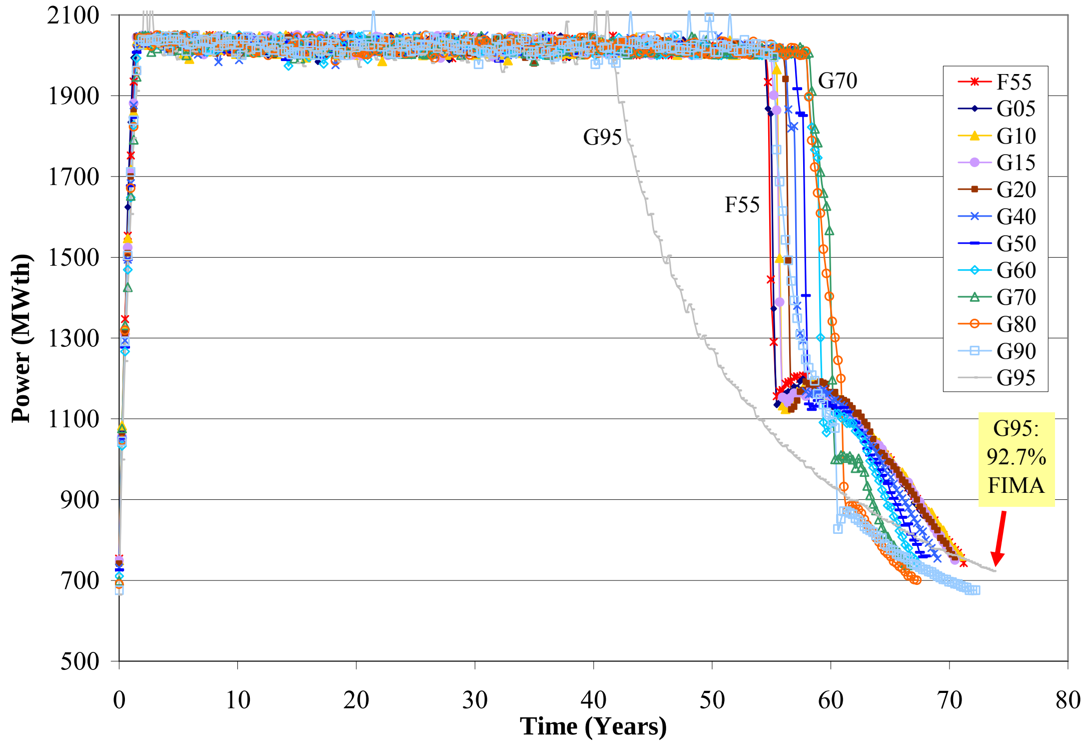

Figure 6. Power Curves for MSFIMS

The next parameters analyzed were the tritium breeding ratio (TBR) and tritium mass curves, found in Figure 7 and Figure 8 respectively. These curves demonstrate the additional tritium production enabled by moderating the neutron spectrum, likely due to less parasitic absorption in the resonance region as well as increasing neutron capture cross sections of lithium as the spectrum becomes more thermal. The TBR and tritium mass curves provide no clear indication of which case is optimal but serve as good references when seeking to understand differences between cases.

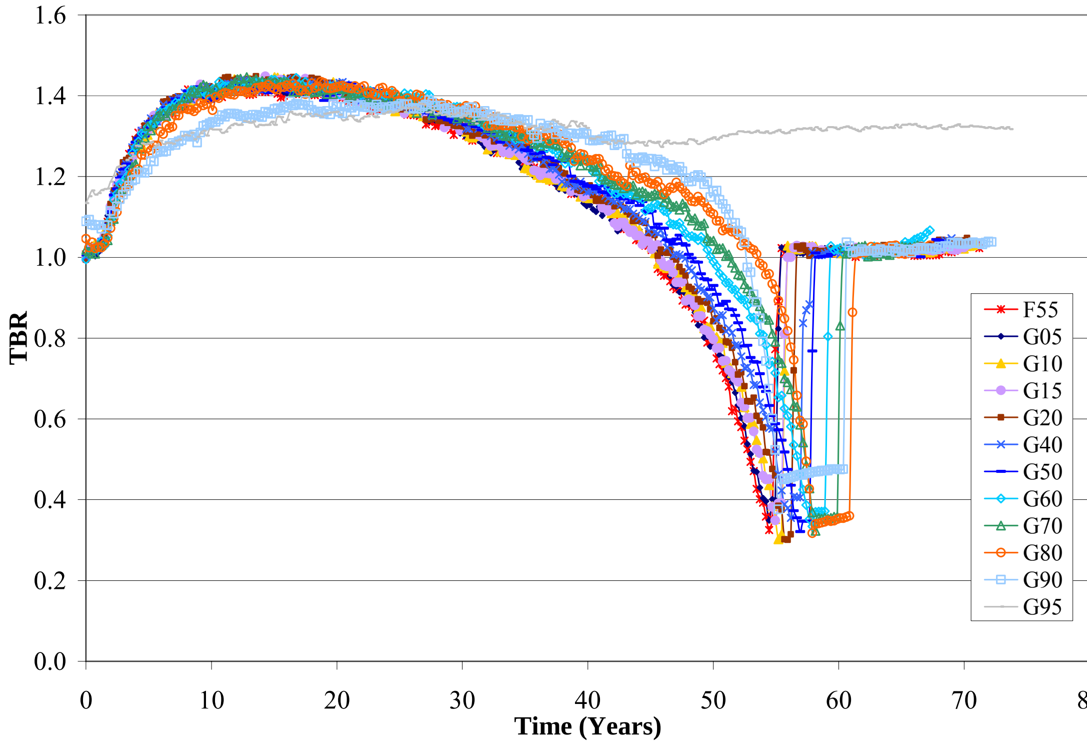

Figure 7. TBR Curves for MSFIMS

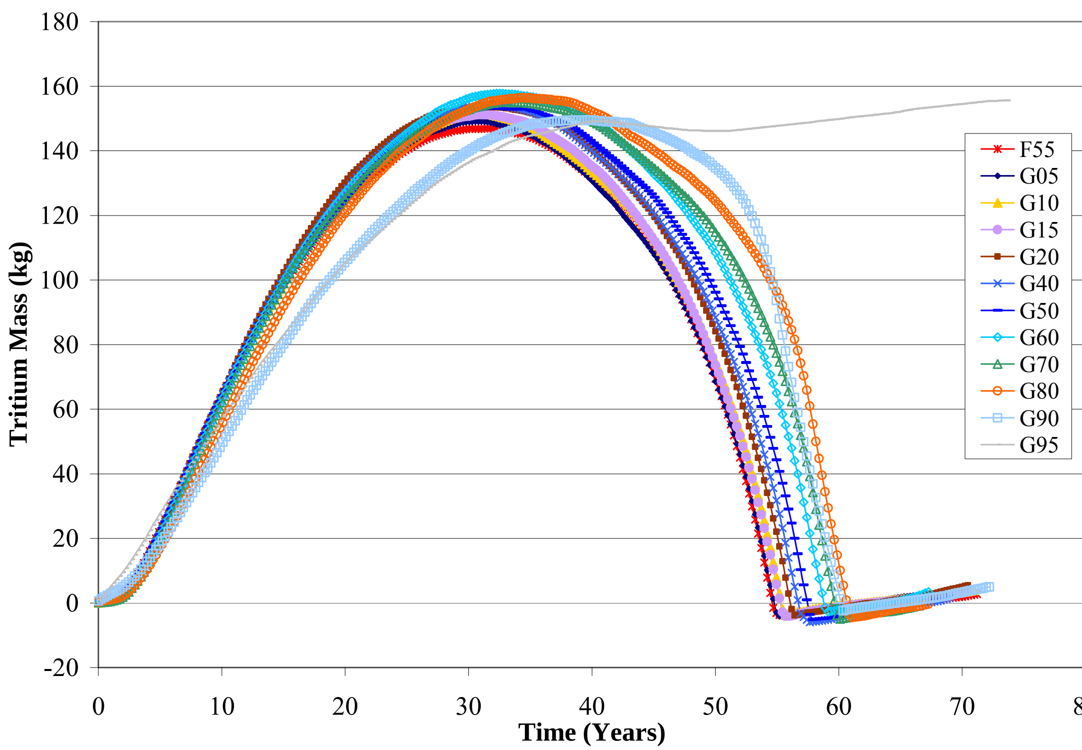

Figure 8. Tritium Mass Curves for MSFIMS

The actinide mass curves for the MSFIMS in Figure 9 provide the key evidence that G70 reached the 99% FIMA burnup point fastest and also show the general trends exhibited by varying moderator percent in the molten salt region of the model. The various cases all burned through the actinide mass at a similar rate through about the first 50 years of operation, but then diverge quickly thereafter. Figure 10 reports the explicit %FIMA burnup curves for the MSFIMS cases after 50 years of operation, showing the same trends and conclusion mentioned above. Figure 11 displays the total system mass and incineration time, defined as the time to reach 99% FIMA burnup, as a function of volume percent of graphite in the molten salt fuel region.

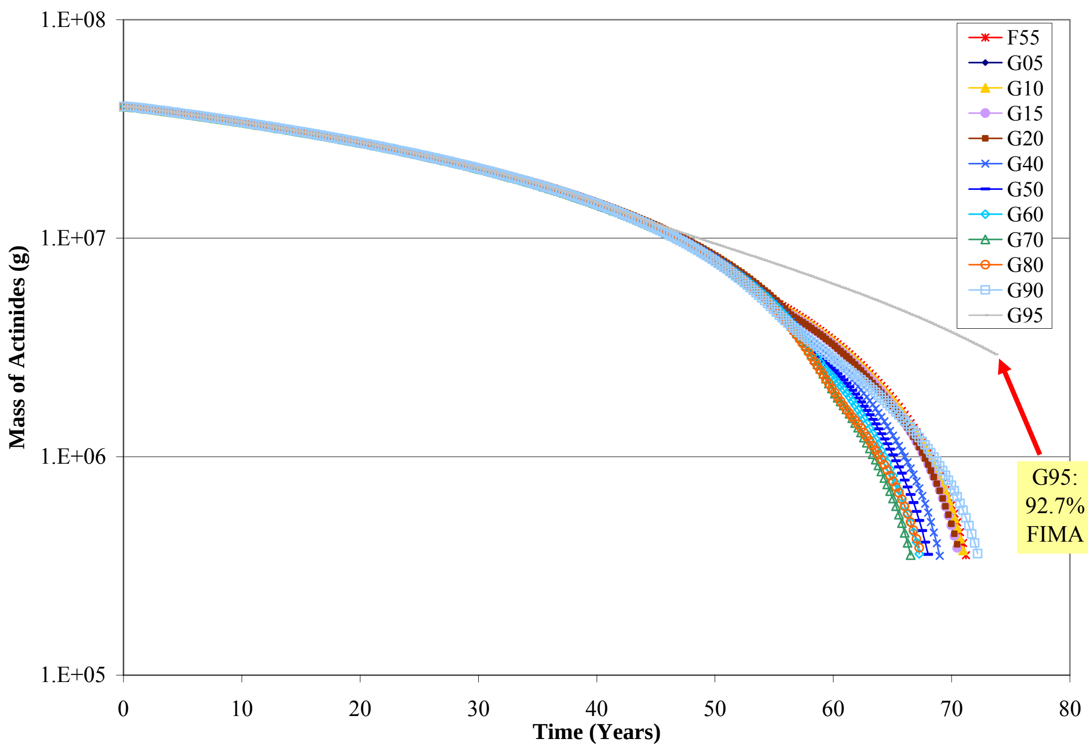

Figure 9. Actinide Mass Curves for MSFIMS

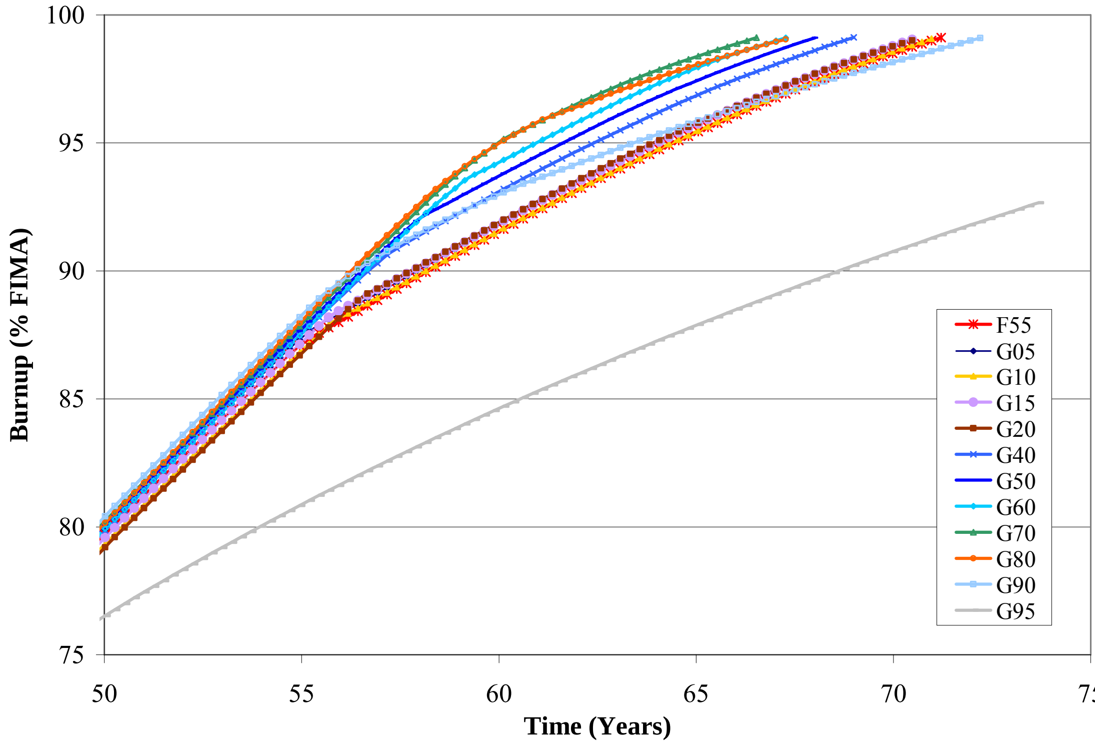

Figure 10. Burnup Curves After 50 Years of Operation for MSFIMS

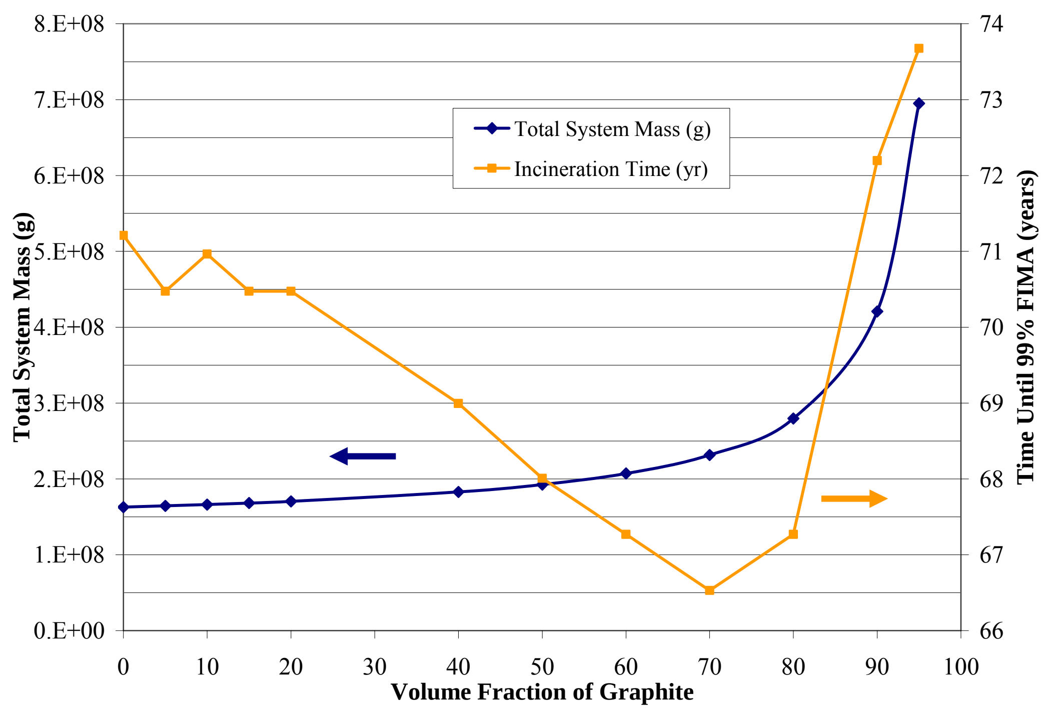

Figure 11. Total System Mass and Incineration Time as a Function of Moderator Volume Percent

Finally, the performance indices defined in Section 3 were calculated for all MSFIMS cases. These performance indices, as shown in Figure 12, show that the molten salt fast ignition system optimizes somewhere between 60 and 80 volume percent of graphite in the molten salt and graphite region if the performance indices are to be maximized. Figure 11, above, demonstrates that this same optimization region is true for minimizing the incineration time.

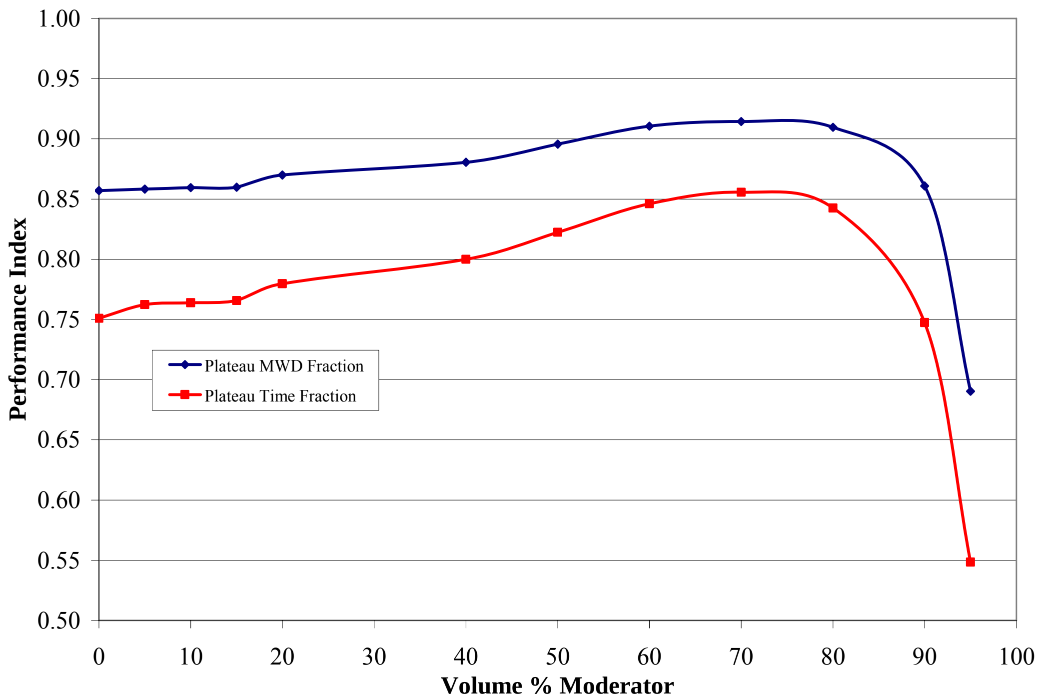

Figure 12. Performance Indices as a Function of Moderator Volume Percent

Using all of the preceding curves and analysis, case G70 was chosen as the reference case for a molten salt blanket fast ignition LIFE engine with a molten salt and graphite region that is 70 vol% graphite and 30 vol% molten salt. It offers the shortest incineration time at about 66.5 years to reach 99% FIMA burnup, has the longest operational time at the peak plateau power level of 2000 MW_{th}, and has the highest performance indices of any case in the moderator study.

A request was made by Ralph Moir to examine the evolution of the molten salt composition during the lifetime of the design. Specifically, it was requested than an analysis be performed to determine the mol percent of various constituents in the molten salt as a function of time. These mol percents were calculated as the percent of total atoms that the atoms of individual elements represented. The number of total atoms did not include fluorine, since all other atoms are bonding with the fluorine, or carbon since almost all of the carbon on the molten salt material would actually be in the form of solid graphite. Carbon atoms may have to be included, however, if graphite shells or pebbles are shown to give off some amount of carbon to the molten salt. Figure 13 displays the results of this analysis.

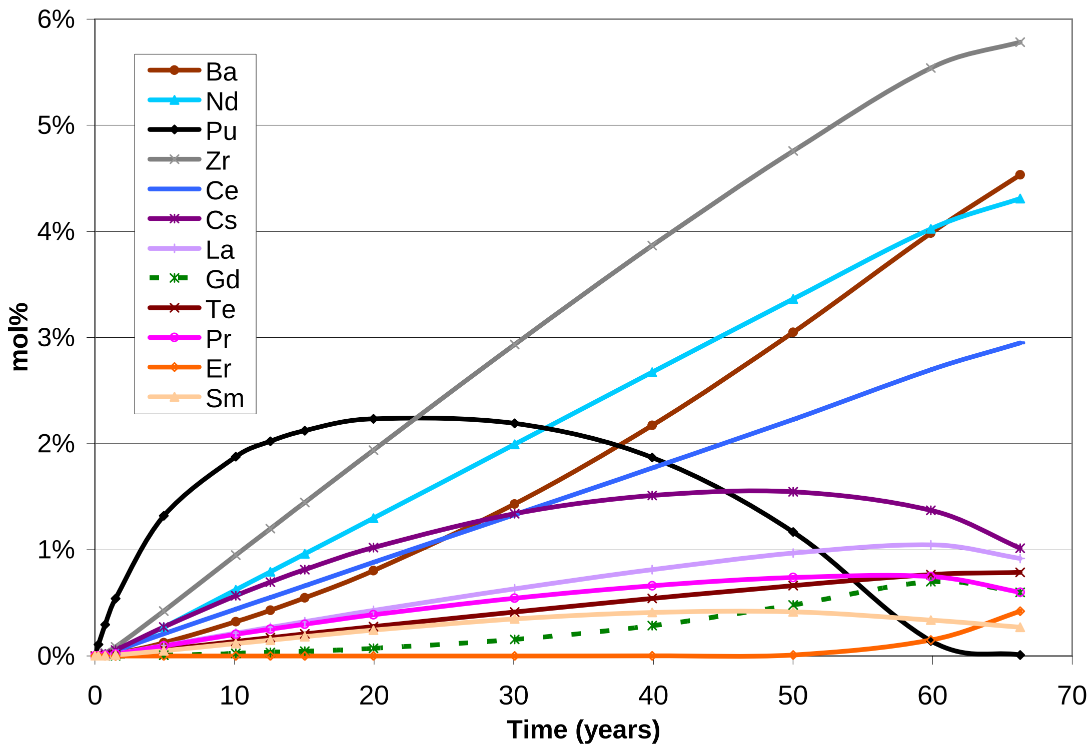

Figure 13. Evolution of Molten Salt Components with Time for Case G70

Based upon initial guidance given that the peak plutonium mol percent should be kept down around 1 mol%, these results indicate that some additional design work needed be performed. This led to another study, which is documented in detail in Section 8.3. It is important to note that more recent guidance in the form of an email from Henry Shaw at LLNL has indicated that the total mol fraction of all trivalent actinides (Pu, Cm, and Am) and rare earth elements out of all cations that form fluoride salts may be a more accurate metric. The solubility limit for this total mixture is a strong function of both temperature and the composition of the balance of the molten salt and is not known at this point for molten salts similar to the ones in this study.

An additional request was made for the isotopic masses of U-232, U-233, Th-232, and the cumulative mass of U-233 and Th-233 as a function of time, along with the time dependent mass ratio of U-232 to U-233. Figure 14 displays the time-dependent isotopic masses of interest for case G70 and Table 4 reports the calculated time-dependent U-232 to U-233 mass ratio for case G70. These results are provided here for information purposes; no conclusions are drawn from them in this analysis.

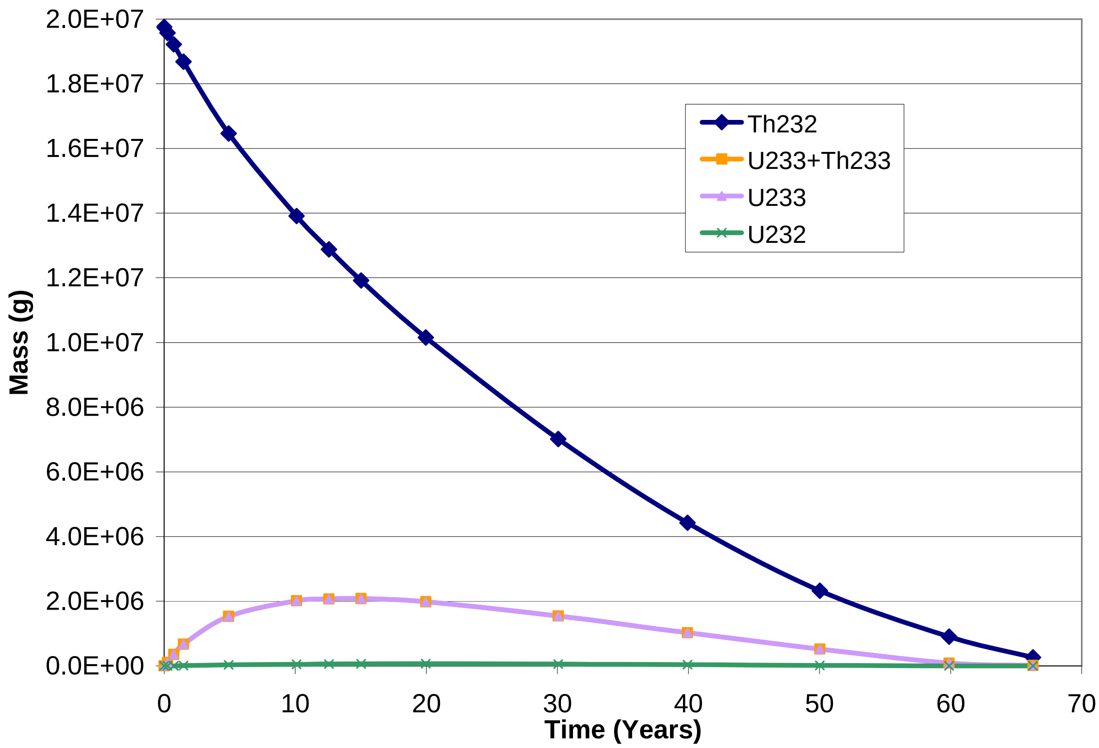

Figure 14. Time-dependent Isotopic Masses for Case G70

Table 4. Time-dependent Mass Ratio of U-232 to U-233 for case G70

Time (Years)	U232/U233 Ratio
0	0.0000
0.25	0.0047
0.74	0.0085
1.48	0.0119
4.93	0.0183
10.1	0.0250
12.57	0.0280
15.03	0.0304
19.96	0.0342
30.06	0.0367
39.92	0.0342
50.02	0.0281
59.88	0.0353
66.28	0.0448

7.2. Comparison to Analogous TRISO Fuel Case

Having established case G70 as the molten salt fast ignition reference case, it is important to compare G70 to an analogous TRISO fuel fast ignition reference case to both better illustrate the features of the molten salt case and also understand any fundamental differences between molten salt and TRISO concepts. For the purposes of this comparison, case “da0” was chosen as the TRISO fuel reference case.

The da0 case used TRISO fuel particles in pebbles as its fuel form, was initially loaded with 40 MT of actinides, had a plateau power of 2000 MW_{th}, and used 500 MW of fusion power from a fast ignition system. The only difference between da0 and G70 in the above parameters is the fuel form used; other differences do exist between da0 and G70 in both fundamental design as well as design performance, however, as described below.

One of the most fundamental differences between da0 and G70 lies in the continuous removal of noble and semi-noble metals and gaseous fission products in G70, as described previously in this report. This difference will be quantified in further detail later, but is important to at least conceptually note at this point.

Another key difference between da0 and G70 surfaces in looking at the masses of various components and the total system mass, as illustrated in Figure 15 below. The beryllium neutron multiplier region mass is identical between the two cases and the mass of the walls is similar, but the other regions exhibit significant differences. The large mass difference in the "Fuel and Moderator" region of the models, when analyzed at the detailed level of isotopic masses, consisted mainly of mass reduction in lithium and carbon and smaller contributions from mass reduction due to elimination of silicon and beryllium in this region. The total system mass from the MCNP model of G70 shows a mass reduction of about 40% relative to the da0 total system mass.

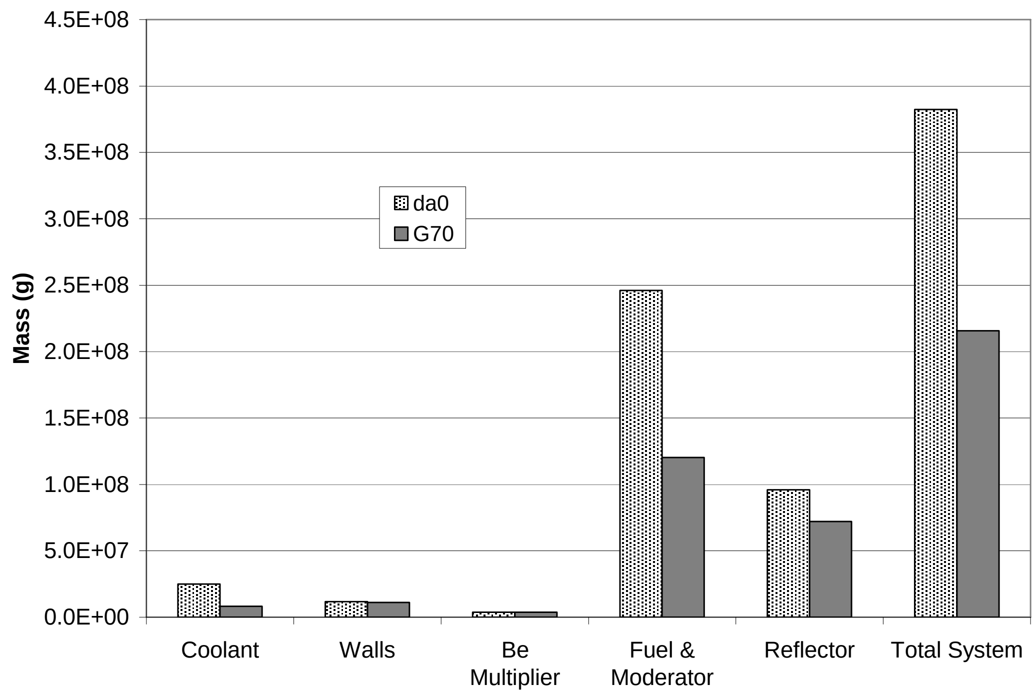

Figure 15. Mass Comparison of Reference Molten Salt and TRISO Cases

A final fundamental difference to note in comparing da0 to the moderator study performed for molten salt fast ignition comes from comparing the moderator-to-fuel ratios on both an atom and mass basis, as displayed in Figure 16.

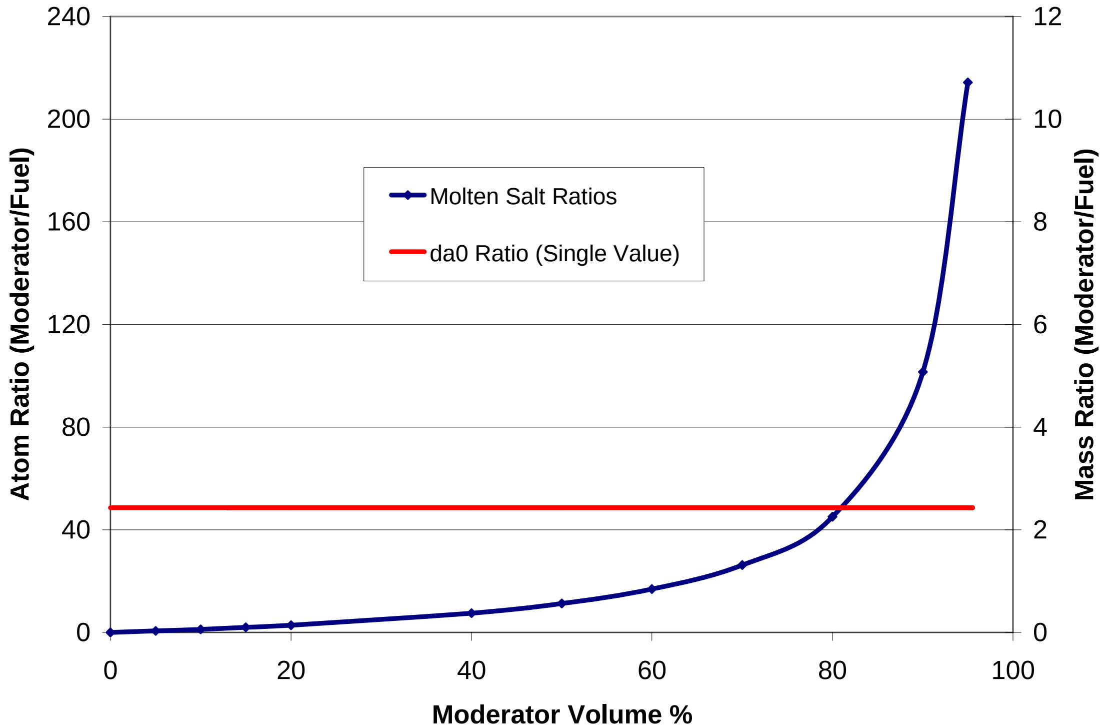

Figure 16. Moderator-to-Fuel Atom and Mass Ratios for Molten Salt and da0

The intersection of the two curves indicates that the molten salt case should optimize somewhere near a moderator volume percent of 80, while the actual results of the study show that a value around 70 volume percent moderator was better. The near proximity of these results provides at least a basic level of independent verification of the molten salt moderator study since the two values are within the same range of values. The differences between da0 and the molten salt curve could be due to various factors including da0 not being fully optimized, less moderation needed in the molten salt case because of less parasitic neutron absorptions due to removing some fission products, or a combination of these two factors and others. The impact of fission product removal will be examined in much further detail later in this report, including looking at different levels of fission product removal.

Having established some of the fundamental differences in the initial models of cases da0 and G70, comparisons of some of the results from the neutronics burnup calculations become the next step in looking at differences and similarities of TRISO and molten salt LIFE engine blanket designs. Figure 17 shows the power curves for da0 and G70.

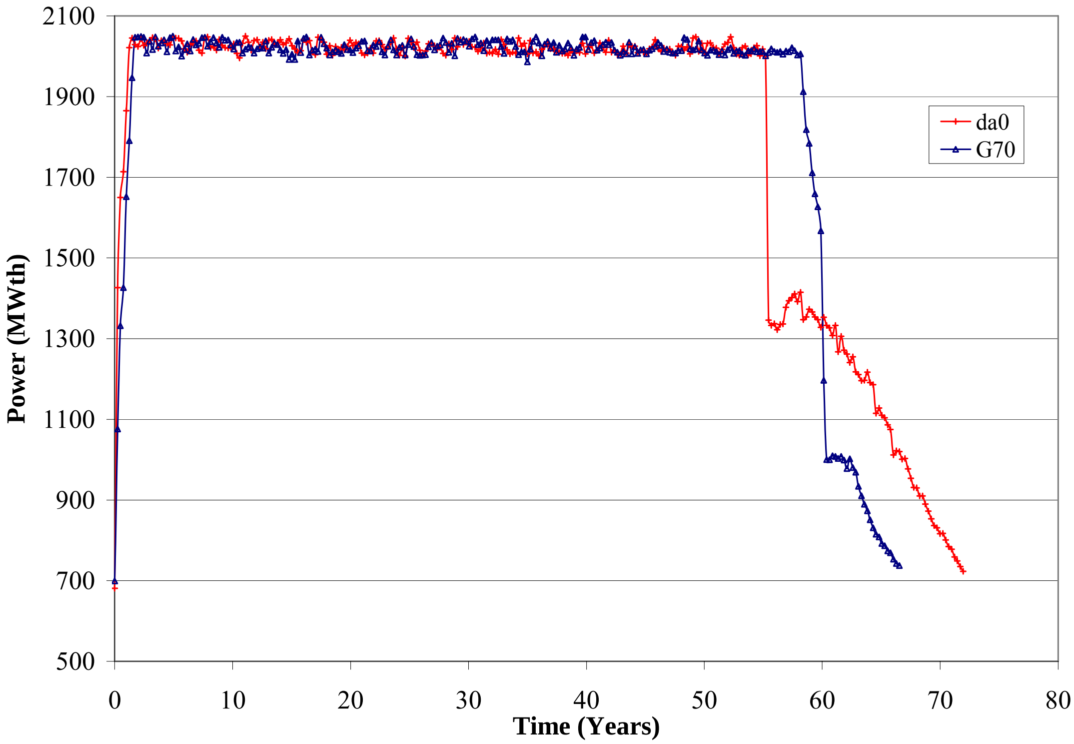

Figure 17. Power Curves for Cases G70 and da0

Based upon the extended plateau region of the G70 power curve, as well as its ending point being before da0, it seems possible that G70 may offer better performance than da0. Other metrics offer further insight into this possibility.

TBR and tritium mass curves from cases da0 and G70 are shown in Figure 18 and Figure 19, respectively. As with the MSFIMS, no direct inference can be made form these curves as to whether either case offers better performance. It should, however, be noted that the peak tritium inventory for G70 (about 155kg) is nearly doubly the peak tritium inventory of da0 (about 86kg). This presents potential issues both in storing and handling such a large tritium inventory as well as a reduction in overall LIFE engine performance since any stored tritium is constantly undergoing radioactive decay and thus tritium storage at high levels can be a waste of neutrons used to produce the material. Either peak tritium inventory represents a very large quantity of tritium and would present many challenges, but the G70 tritium inventory would be more challenging. It is worth noting that the high tritium production in G70 indicates that the system naturally wants to come up to a higher power level.

It important to note that a sharp change in the TBR and tritium mass curves for case G70 can be seen at a time of around 50 years. This change, present not only here but in various other TBR and tritium mass curves in this report, stems from approaching the transition from Phase 2 to Phase 3 operation as defined in Section 1. The system in case G70 has consumed most of the fissile material in the molten salt and allows the tritium breeding to decrease in an effort to take neutrons that were being used to breed tritium and instead use them to maintain plateau power by converting more fertile material to fissile material.

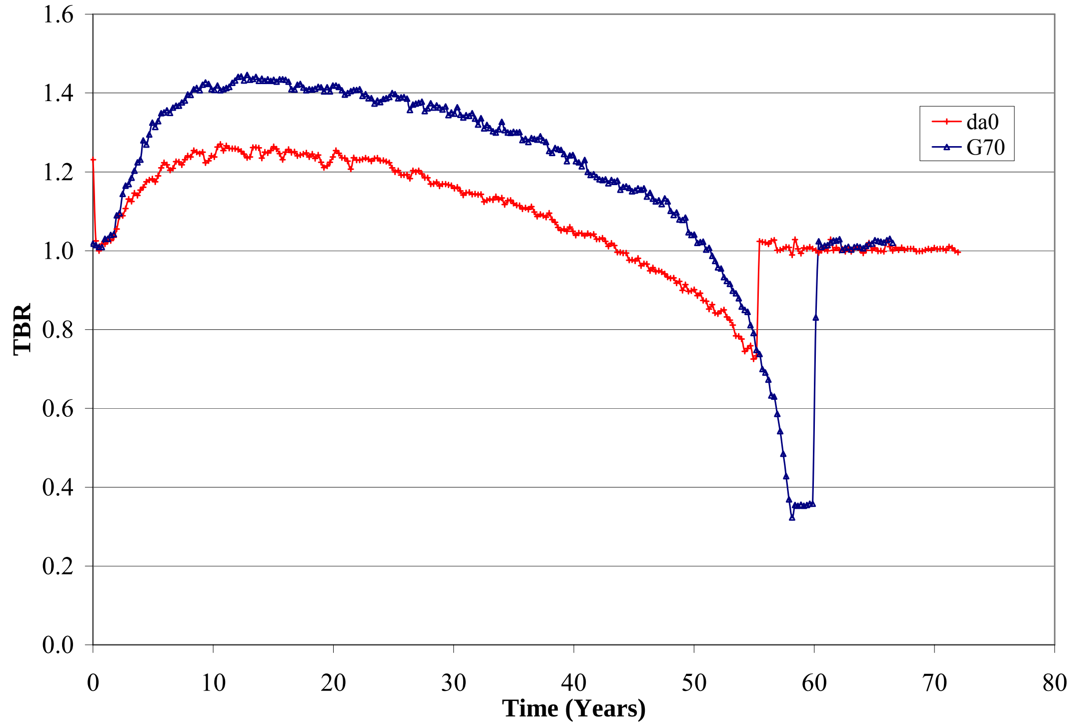

Figure 18. Comparison of G70 and da0 TBR Curves

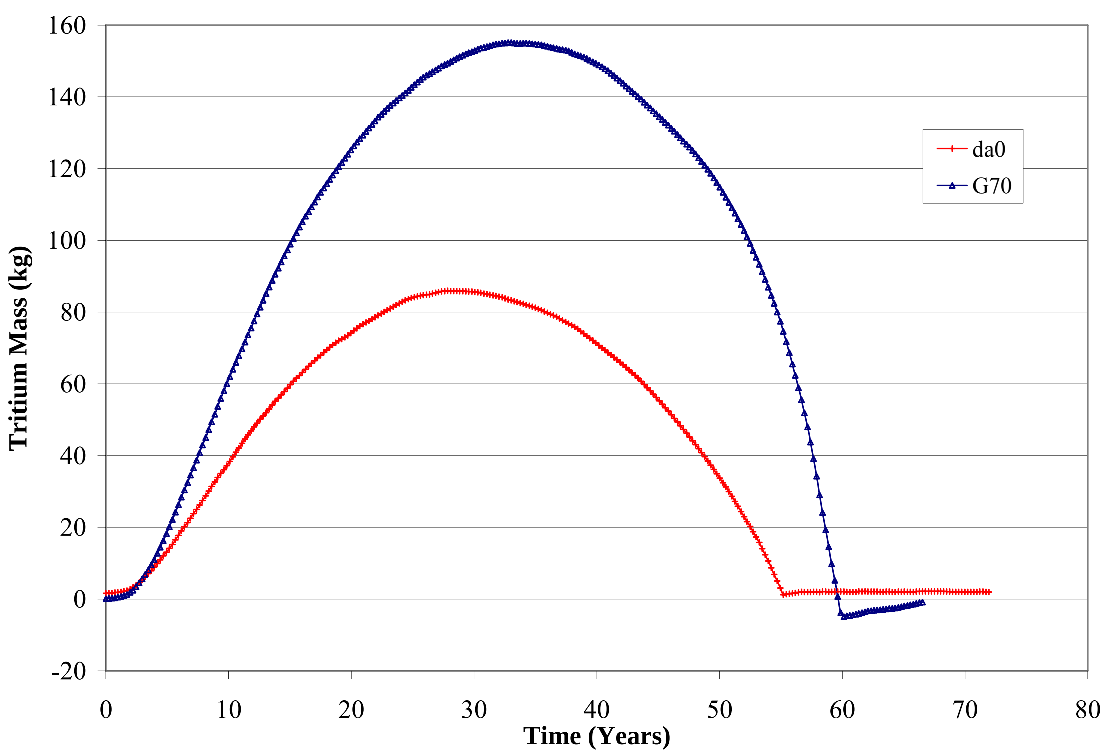

Figure 19. Tritium Mass Curves for Cases G70 and da0

The actinide mass curves for cases G70 and da0, shown below in Figure 20, provide some assurance that both cases are reasonable since they display similar trends through independent calculations and analyses. G70 incinerates the 40 metric tons of actinides in about 5.5 less years than da0 requires, which is a small but significant advantage.

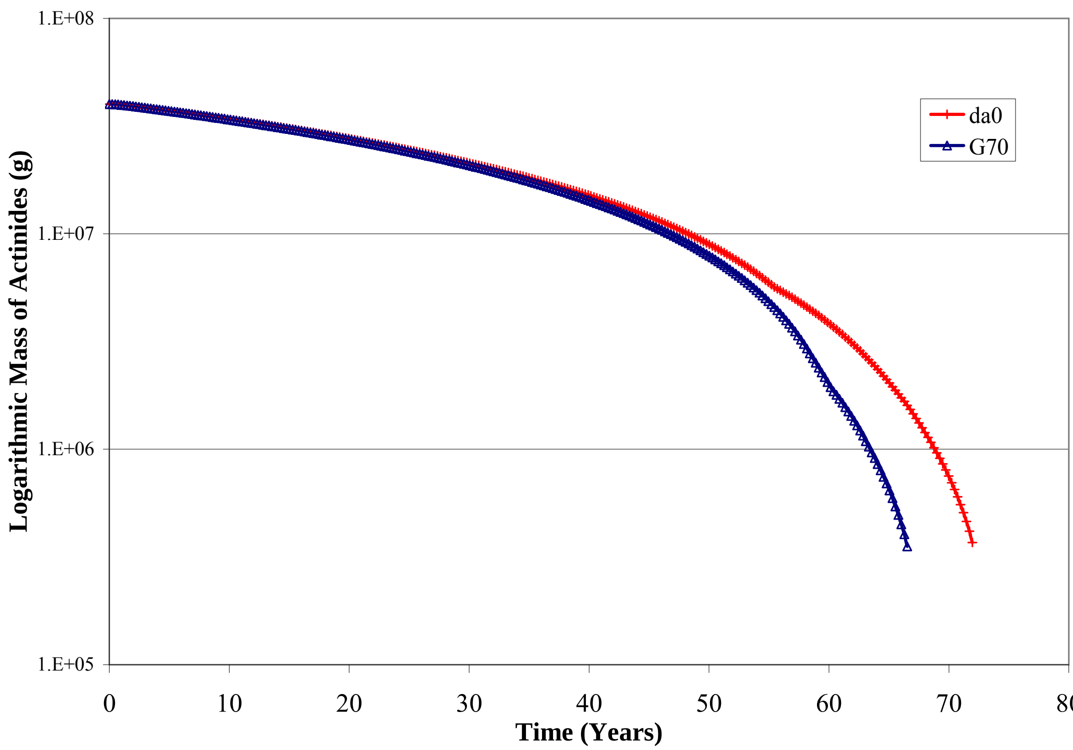

Figure 20. Actinide Mass Curves for Cases G70 and da0

As with the MSFIMS, the performance indices defined in Section 4 allow for better clarity in seeing and comparing results. Figure 21 shows the performance indices for da0 and G70 and indicates that G70 offers potentially better performance than da0.

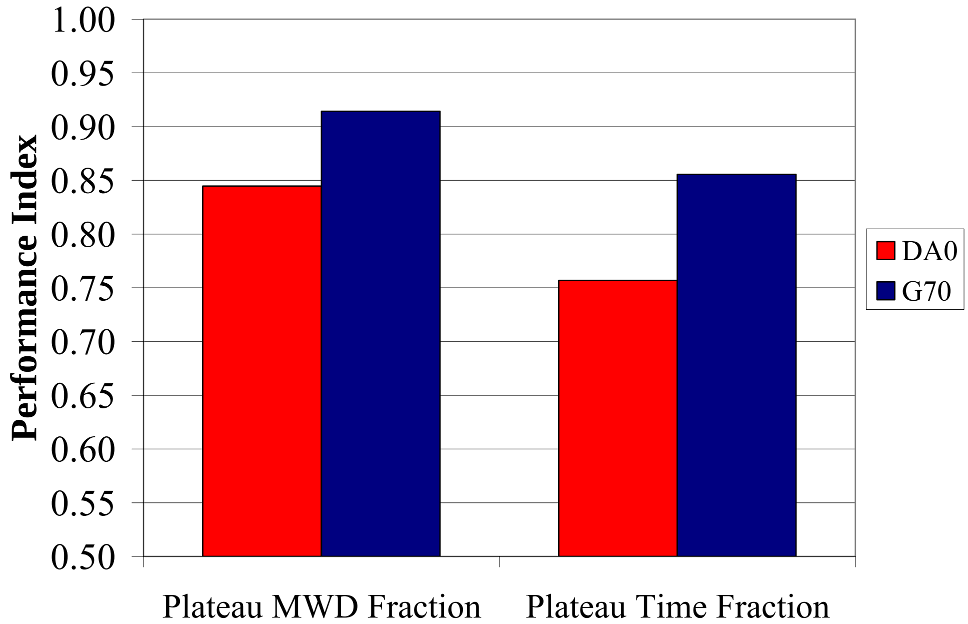

Figure 21. Performance Indices for Cases G70 and da0

The fast ignition reference cases molten salt (case G70) and TRISO (case da0) based blanket designs contain some clear fundamental differences in design approaches that translate into differences in masses and performance. G70 appears to offer some performance benefits relative to da0 in the form of shorter incineration time and higher performance indices, but more detailed design work needs to be performed.

8. SCOPING STUDIES FOR ADDITIONAL CONCEPTS

This section contains descriptions of, and results from, scoping studies that were performed to investigate additional concepts that were identified as being of interest during the Molten Salt Fast Ignition Moderator Study.

8.1. Continuous Removal of Miscellaneous Additional Gases

8.1.1. Purpose

The reference case identified in the MSFIMS, case G70, did not include the continuous removal of miscellaneous gasses that were shown to form during operation. These gases included hydrogen (H-1 and H-2 species), helium (He-3 and He-4 species), oxygen (O-16 and O-17 species), and neon (Ne-20). Through emails with Jeff Latkowski at LLNL, it was agreed upon that these gases would either naturally dissociate from the salt or would be actively removed. This scoping study, identified as case RM1, examined the impact of a reference case variation removing these gases. The effects are hypothesized to be trivial but were explicitly investigated and documented here for verification purposes.

8.1.2. Methods and Model Description

Case RM1 used identical input files to the reference case (case G70) other than specifying that H-1, H-2, He-3, He-4, O-16, O-17, and Ne-20 should be removed during the TBR code's preprocessing routines.

8.1.3. Results

The results from this study proved, as was expected, that the differences due to removing the list of gases described above were indeed trivial. The power curves for G70 and RM1 are shown in Figure 22 and the tritium mass curves for G70 and RM1 are shown in Figure 23; both of these figures show that the differences in results are essentially negligible since they are within the stochastic uncertainties of the MCNP results.

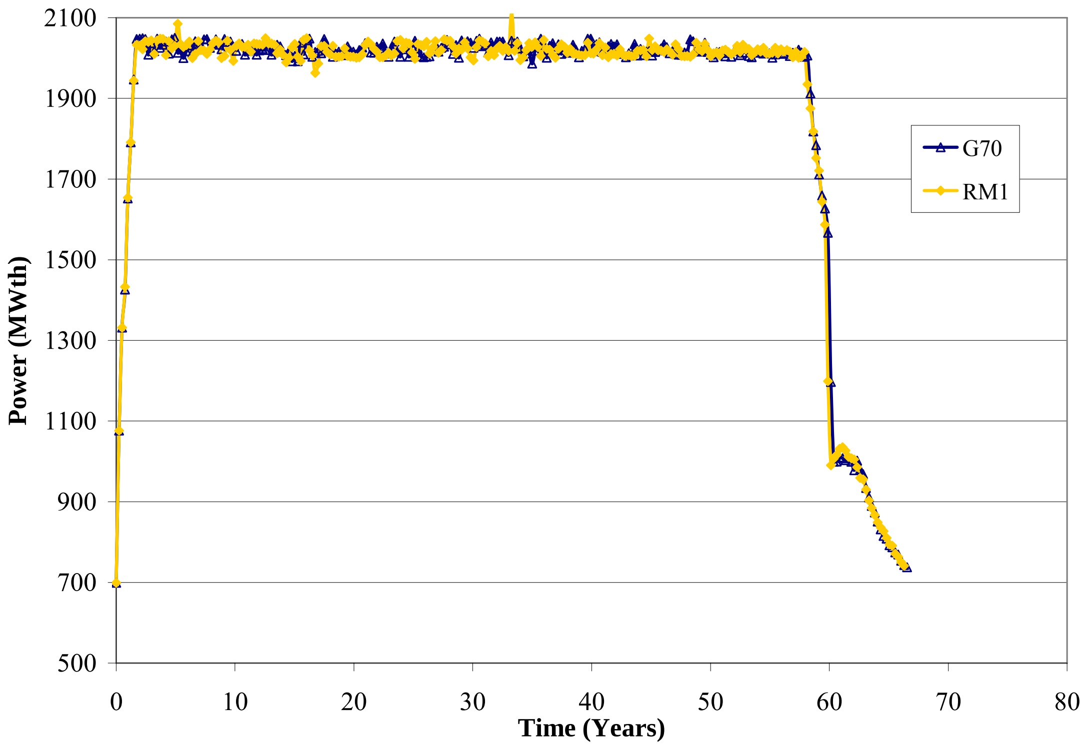

Figure 22. Comparison of G70 and RM1 Power Curves

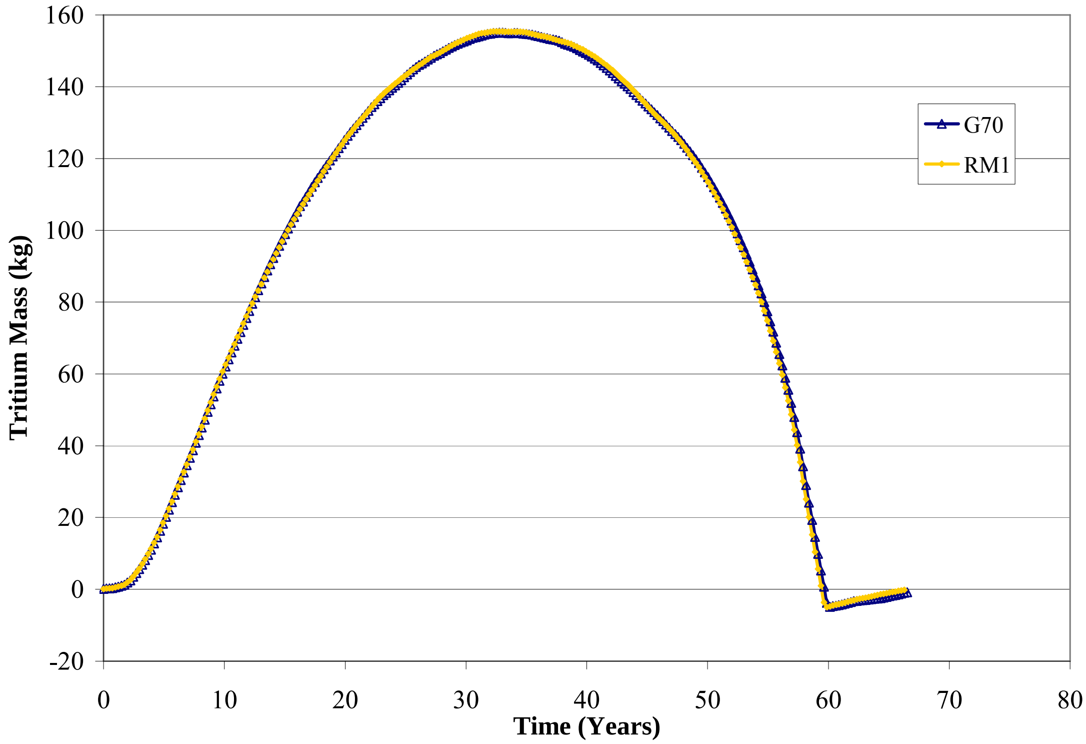

Figure 23. Comparison of G70 and RM1 Tritium Mass Curves

The comparison of G70 to RM1 shows that G70 is an acceptable fast ignition reference case for a molten salt blanket design.

8.2. Continuous Removal of Rare Earth Element Fission Products

8.2.1. Purpose

This scoping study, identified as case RM2, originated from an email request from Ralph Moir to investigate the effects of removing all of the rare earth element fission products in addition to the noble and semi-noble metals and gaseous fission products already continuously removed in the reference case (case G70). Solubility concerns regarding the rare earths and other elements they affect were given as the reason for removal of rare earth elements in addition to the other fission products removed in G70.

Subsequent to the initial scoping study performed through case RM2, another case identified as RM3 was run. RM3 sought to convert the excess tritium mass at 99% FIMA found during the RM2 study into a higher plateau power and thus enable a more compelling design. RM3 changed the plateau power level from the 2000 MW_{th} used in the MSFIMS and case RM2 to 2200 MW_{th}.

8.2.2. Methods and Model Description

Case RM2 used the same input files as case G70 other than some modifications to the "tbr.inp" file. This file was first modified to specify that the miscellaneous gases mentioned in case RM1 above (H-1, H-2, He-3, He-4, O-16, O-17, and Ne-20) should be removed during the TBR code's preprocessing routines. In addition, a more extensive list of continuous removal was added to include all of the elements of the rare earth elements (La, Ce, Pr, Nd, Pm, Sm, Eu, Gd, Tb, Dy, Ho, and Er). In an effort to provide a clear understanding of what it means to continuously remove rare earth elements in addition to continuous removal of noble and semi-noble metals and gaseous fission products, Table 5 provides a list of the elements that are actually present at 99% FIMA burnup for case RM2. All other elements have either not been produced or have been actively removed. It should be noted that it has been suggested that iodine, cesium, and other elements still included would likely be selected for removal as well.

Table 5. List of Elements Still Present at 99% FIMA for Case RM2

Atomic #	Element
3	Li
4	Be
9	F
34	Se
35	Br
37	Rb
38	Sr
39	Y
40	Zr
52	Te
53	I
55	Cs
56	Ba
82	Pb
90	Th
91	Pa
92	U
93	Np
94	Pu
95	Am
96	Cm
97	Bk
98	Cf
99	Es

Case RM3 used the same input files as case RM2 other than specifying that the desired plateau power level should be 2200 MW_{th} rather than 2000 MW_{th}.

8.2.3. Results

The results of case RM2 indicate a more attractive case than case G70, which was expected since the rare earth elements continuously removed during case RM2 are almost all elements with high neutron absorption cross sections. The power curves of G70 and RM2, found in Figure 24, show a slightly longer plateau and shorter incineration time for RM2 than for G70.

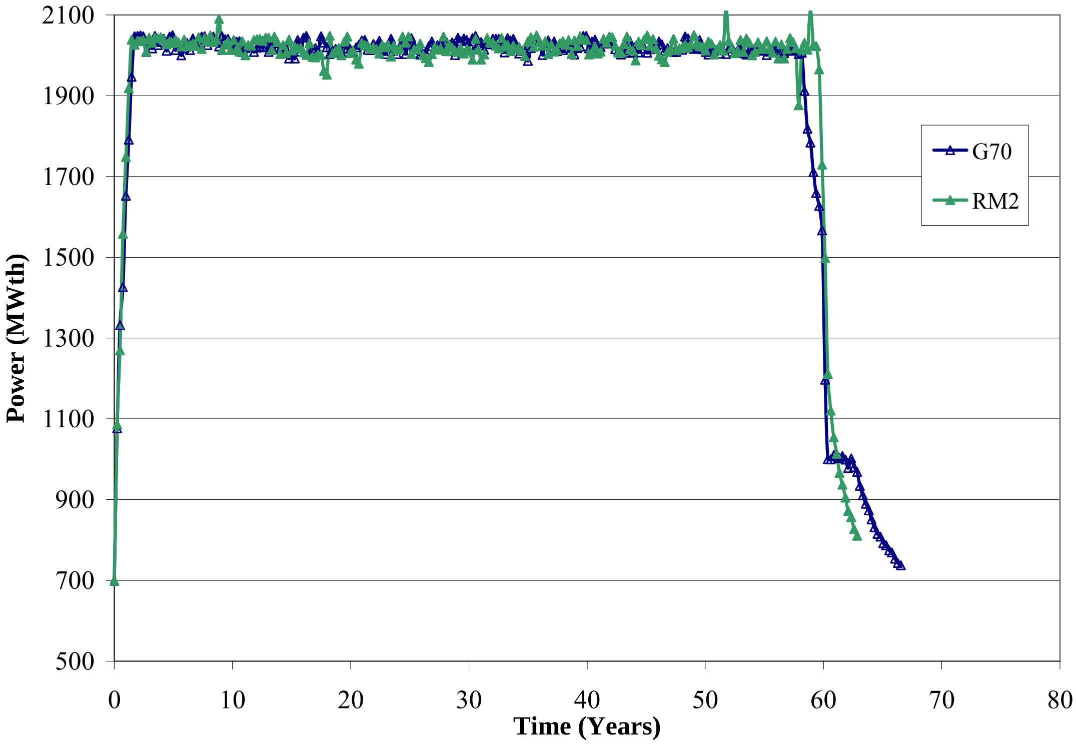

Figure 24. Power Curves for Cases G70 and RM2

Figure 25, which displays the TBR curves for G70 and RM2, indicates that the removal of rare earth elements during RM2 leads to a much higher TBR. The likely reason hypothesized for this is that the lithium in case RM2 is being used to hold the power down to the plateau power level of 2000 MW_{th} even though it naturally wants to rise more than in G70 due to the absence of parasitic neutron absorptions from rare earth elements. As mentioned previous, the use of lithium in this manner is similar to the use of burnable poisons or other reactivity control mechanisms in nuclear fission reactors, with the added benefit that tritium is produced and can be used as fusion fuel at some later point during operation or for another application (e.g., a different LIFE engine).

Figure 26 displays the resulting tritium mass curves and indicates that RM2 reaches 99% FIMA burnup with about 86kg of excess tritium still available. This tritium could again be used for other applications such as a different LIFE engine, or the system could be redesigned to run at a higher plateau power level as in case RM3. The sharp decrease seen in the TBR and tritium mass curves for case RM2 at around 60 years stems from the high burnup level of the system leaving very little fissile or fissionable material available, as described previously in Section 7.2.

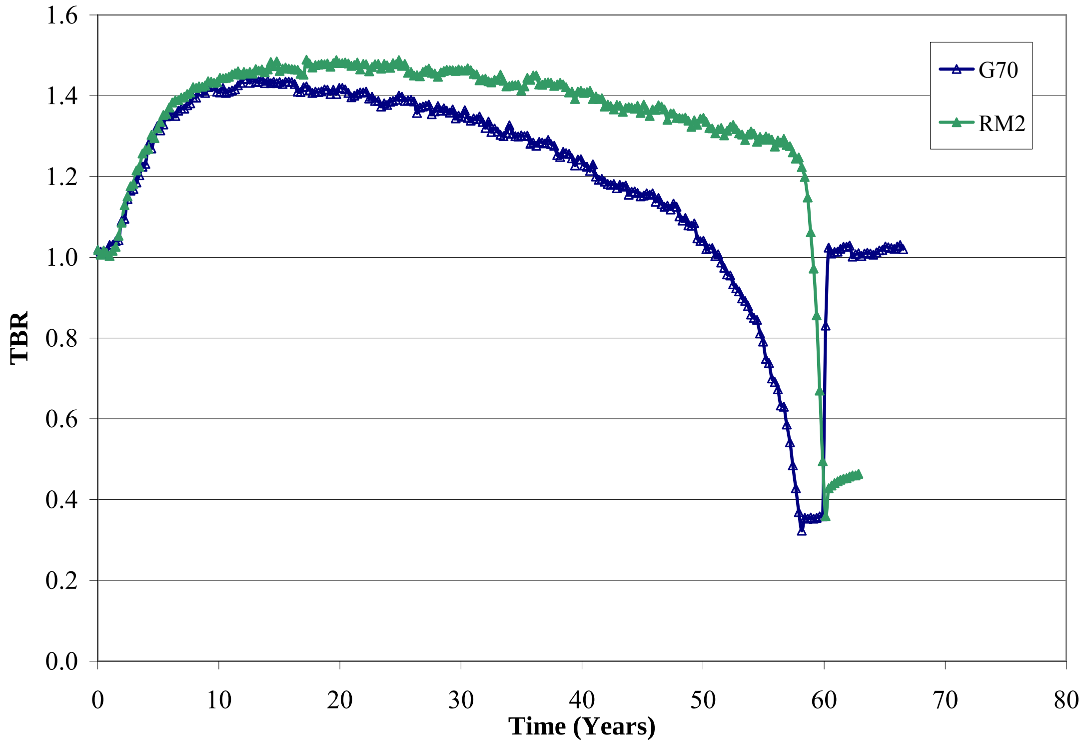

Figure 25. TBR Curves for Cases G70 and RM2

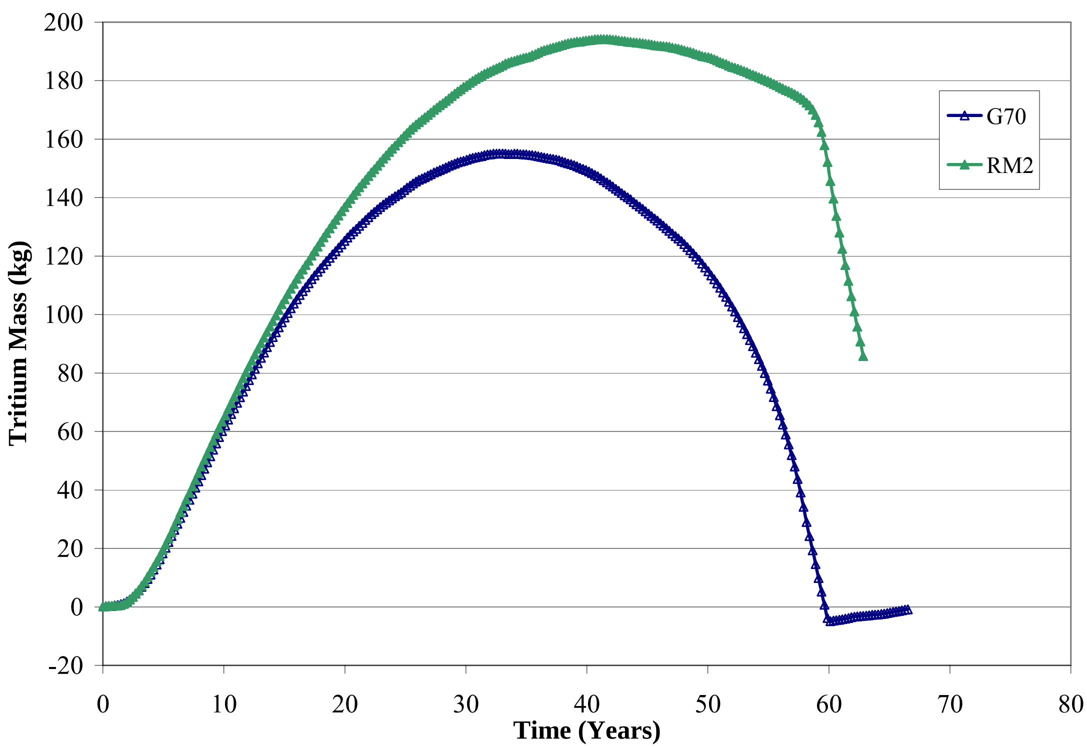

Figure 26. Tritium Mass Curves for Cases G70 and RM2

The actinide mass curves for cases G70 and RM2, found below in Figure 27, indicate an incineration time for RM2 of about 63 years. This incineration time is about 3.5 years shorter than G70.

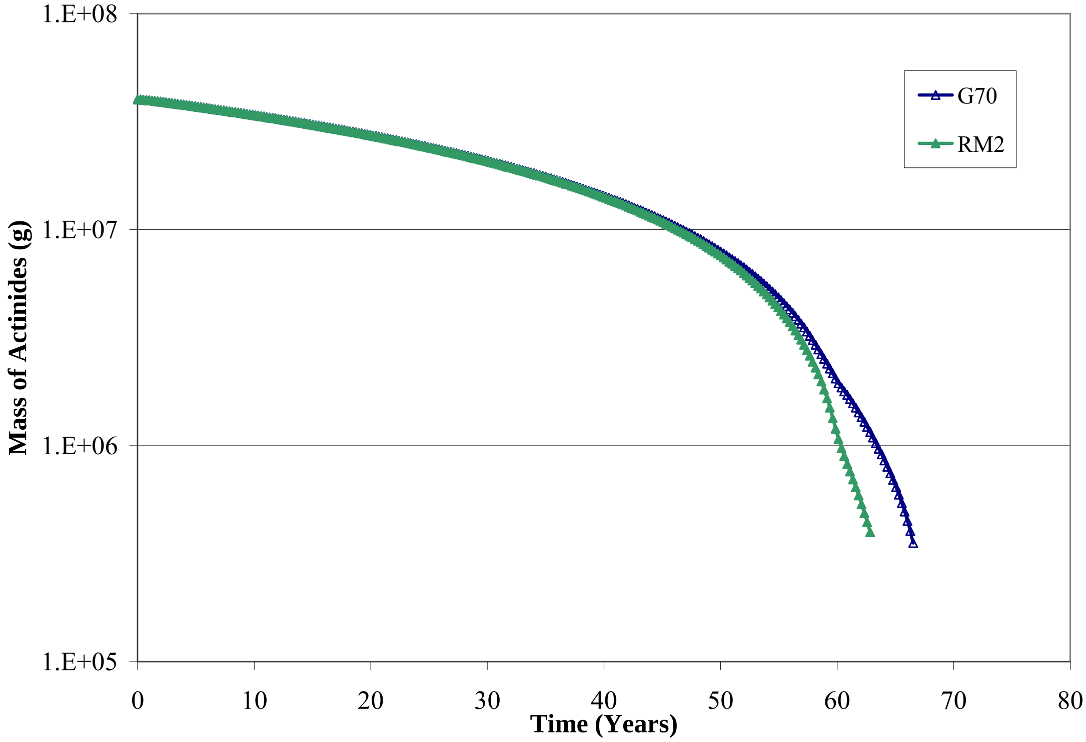

Figure 27. Actinide Mass Curves for Cases G70 and RM2

As with G70, analyses were performed to calculate the time-dependent molten salt composition at an elemental level as well as time-dependent isotopic masses. The time-dependent composition is shown as mol percentages of various elements in Figure 28, Figure 29 shows the time-dependent isotopic masses of interest, and Figure 29 reports the calculated U-232 to U-233 mass ratios. The peak plutonium remained at about 2.25 mol% and still occurred around 20 years after the start of operation. The peak U-233 loading remained around 2000 kg.

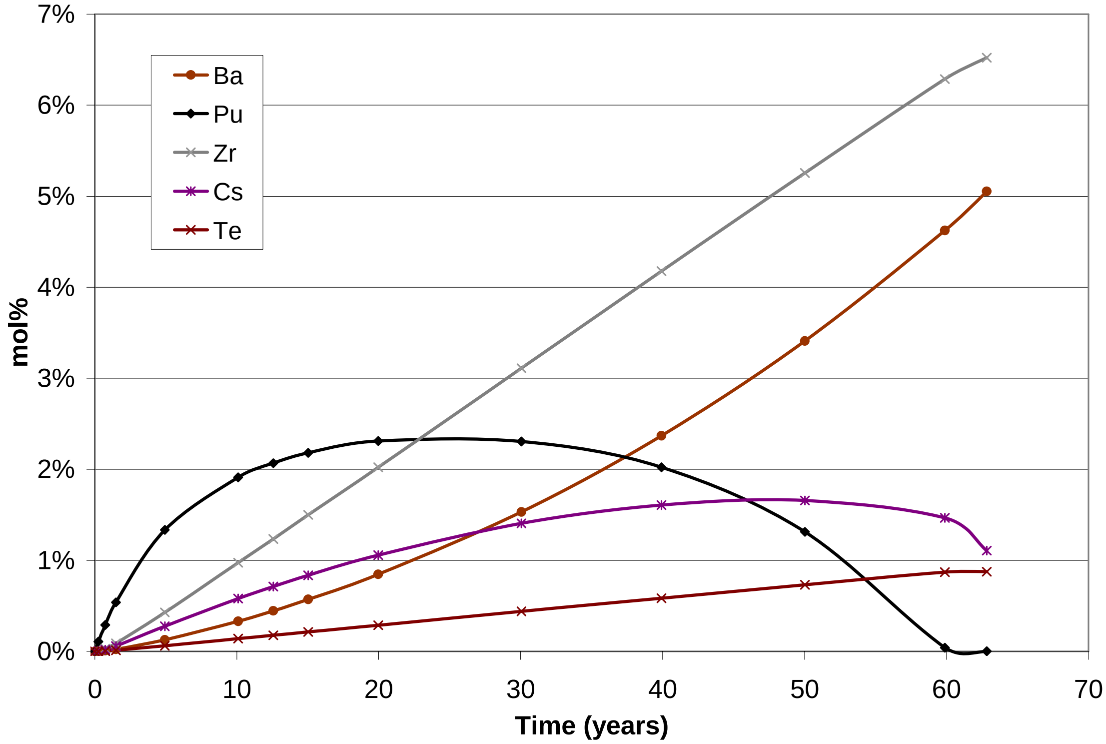

Figure 28. Evolution of Molten Salt Components with Time for Case G70

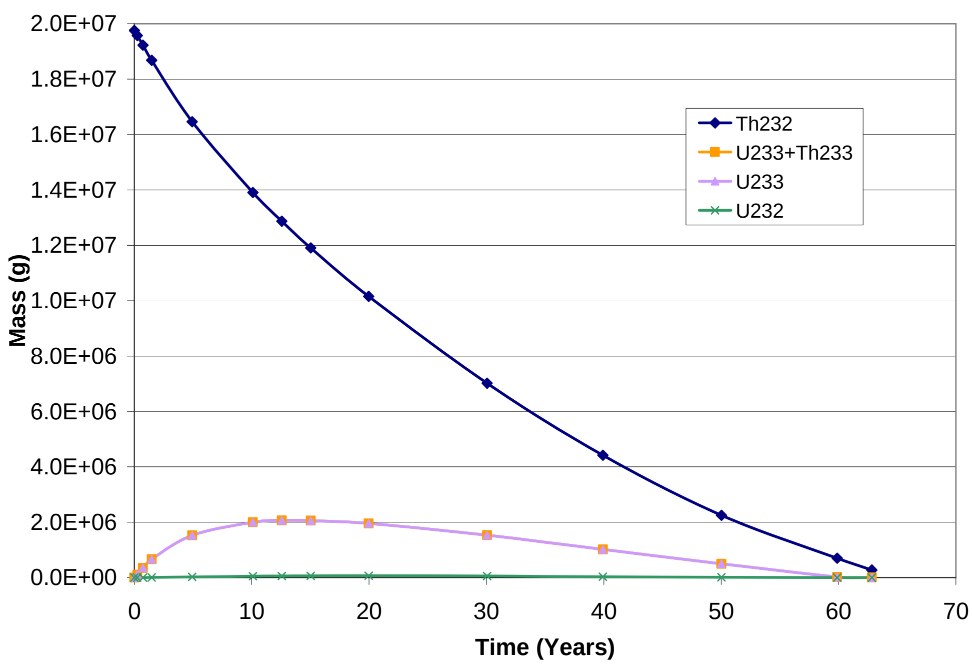

Figure 29. Time-dependent Isotopic Masses for Case RM2

Table 6. Time-Dependent Mass Ratio of U-232 to U-233 for Case RM2

Time (Years)	U232/U233 Mass Ratio
0	n/a
0.25	0.0047
0.74	0.0084
1.48	0.0122
4.93	0.0181
10.1	0.0248
12.57	0.0277
15.03	0.0305
19.96	0.0343
30.06	0.0369
39.92	0.0332
50.02	0.0264
59.88	0.0470
62.83	0.0209

The results of RM2 showed significant improvements over G70 and, as described previously, opened a new study to investigate raising the plateau power from 2000 MW_{th} to 2200 MW_{th}. This subsequent study, case RM3, also yielded promising results. The plots below provide comparisons between cases G70, RM2, and RM3.

Figure 30 provides the power curves for all three cases. The 2200 MW_{th} plateau power of RM3 sharply differs from the 2000 MW_{th} of G70 and RM2 in terms of the power it occurs at, and RM3 ends before the other two due to reaching 99% FIMA faster. The plateau is shorter for RM3 but it has much less non-plateau operation than G70 or RM2, as will be shown later in comparing performance indices. This result is reasonable since actinide mass is conserved and must be fissioned, so the same total energy output is expected between the cases. Since RM3 is run at a higher power, it needs less time to reach the same integrated energy.

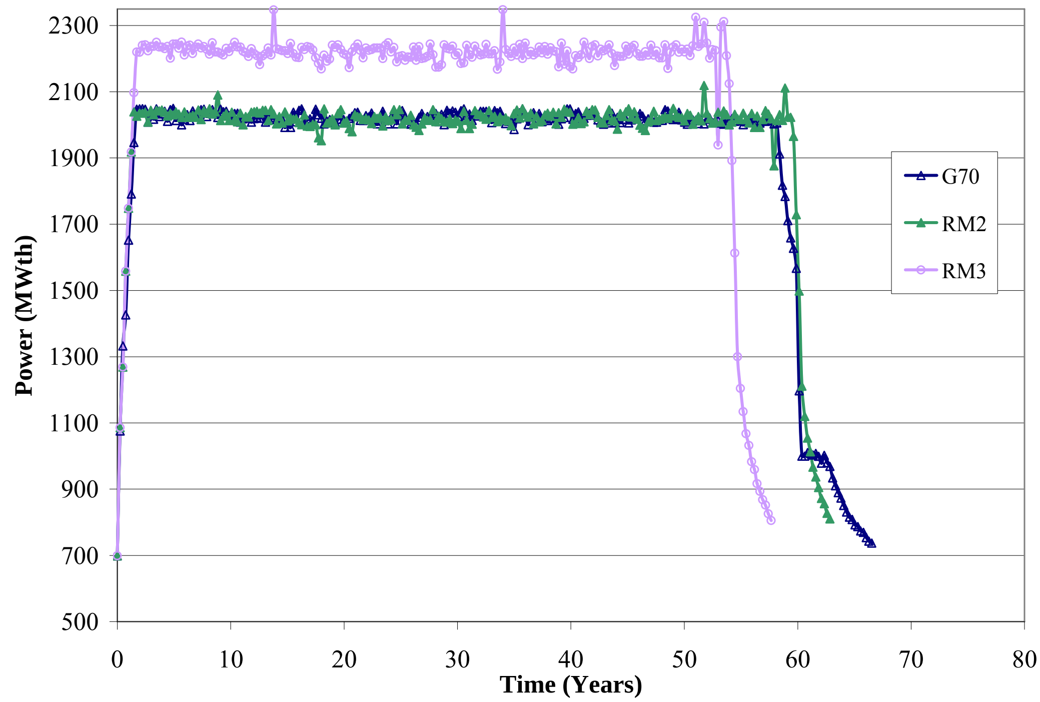

Figure 30. Power Curves for Cases G70, RM2, and RM3

The tritium mass curves in Figure 31 indicate that there is still about 58kg of excess tritium at 99% FIMA in case RM3. This represents a reduction in excess tritium of 5kg compared to RM2, but indicates that the plateau power level could be pushed higher still. A sharp discontinuity can be seen at around 55 years for case RM3 and at around 60 years for case RM2. This is due to a sharp reduction in the TBR, shown in , as more neutrons are needed to create fission events and convert fertile material to fissile or fissionable material and thus less neutrons are available for tritium production.

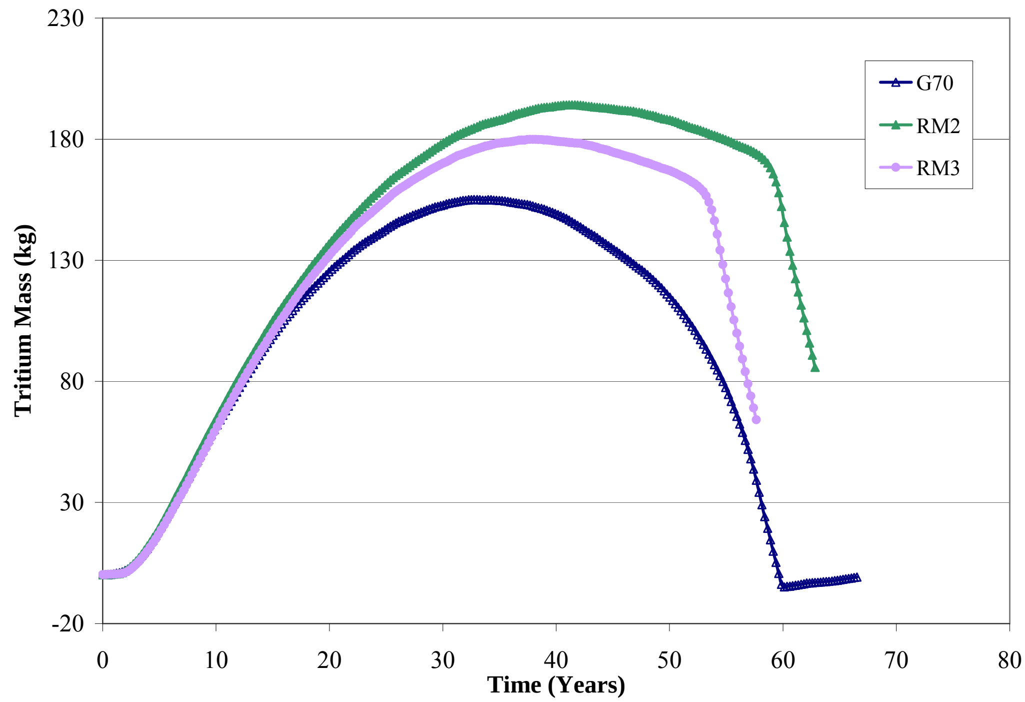

Figure 31. Tritium Mass Curves for Cases G70, RM2, and RM3

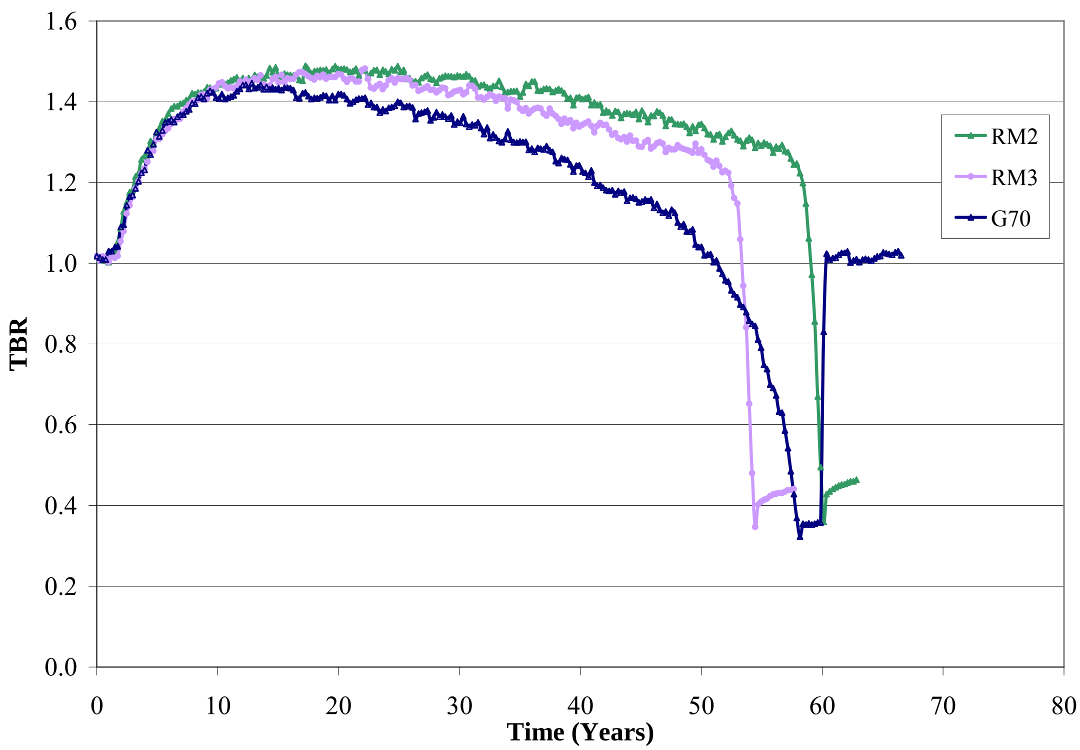

Figure 32. TBR Curves for Cases G70, RM2, and RM3

The actinide mass curves shown in Figure 33 confirm the shorter incineration time conclusions drawn from the power curves. The incineration time for RM3 is approximately 58 years, which is about 5 years shorter than RM2 and about 7.5 or 8 years shorter than G70. This is a substantial reduction in incineration time.

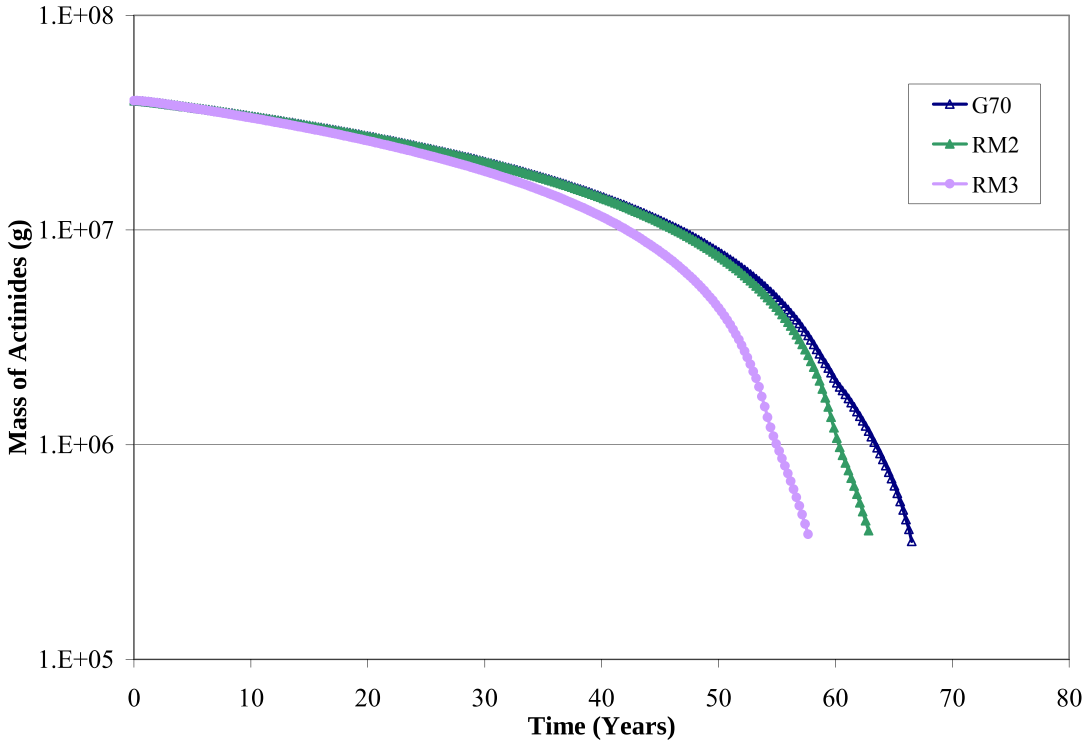

Figure 33. Actinide Mass Curves for Cases G70, RM2, and RM3

Changing the parameters of RM2 to have a 2200 MW_{th} plateau power in RM3 yielded an incineration time 5 years shorter than RM2 and 7.5 or 8 years shorter than G70, allowed for a higher total thermal power for the same fusion power input, and appeared to offer better overall performance. This overall performance was quantified for cases G70, RM2, and RM3 in the form of performance indices which are displayed in Figure 34. RM2 and RM3 both have higher performance indices than G70, and appear to be within resolution limits of each other due to both stochastic error and the resolution offered by the 90 day time steps used for the burnup calculations.

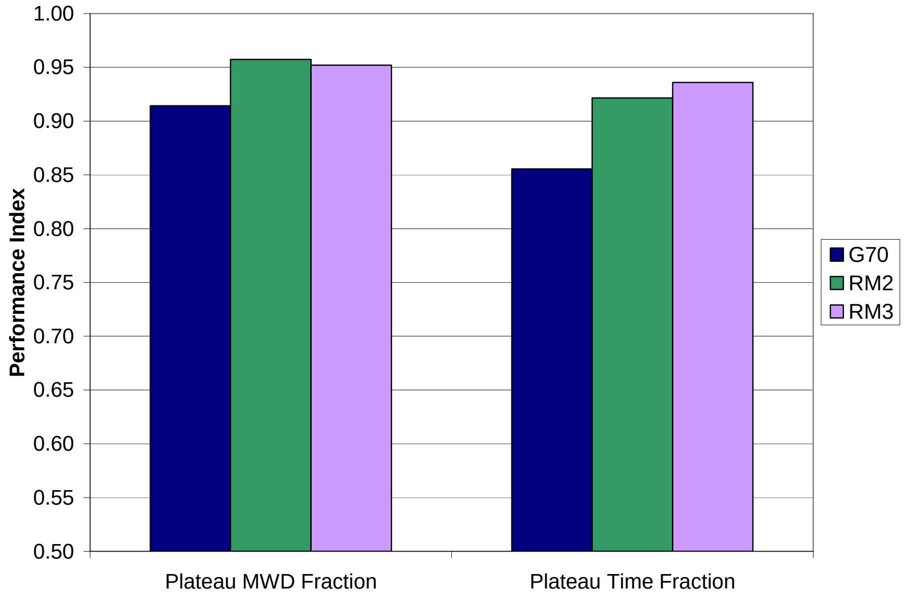

Figure 34. Performance Indices for Cases G70, RM2, and RM3

Based upon all of the previous analyses and figures, case RM3 seems to be the most compelling case among G70, RM2, and RM3. The results for RM3 indicate an incineration time of just under 58 years at a plateau power level of 2200 MW_{th} and very high performance indices. The tritium mass curves, as discussed previously, indicate that the plateau power level could be pushed even higher to leverage the excess tritium at 99% FIMA and thus yield a design with a higher plateau power level and shorter incineration time.

8.3. Alternate Molten Salt Composition

8.3.1. Purpose

The composition of the molten salt used as a basis for nearly all of the analyses in this report stems from initial work performed by Peter Song at LLNL; many alternate molten salt compositions could be investigated with different mixtures of initial loading of plutonium, enriched or depleted uranium, thorium, actinides, or spent nuclear fuel from commercial light water reactors.

This study investigates one possible alternate initial molten salt composition for a fast ignition case. Based upon a conversation with Jeff Latkowski at LLNL, this study changed the reference case 12 mol% UF_4 and 12 mol% ThF_4 molten salt to a 6 mol% UF_4 and 18 mol% ThF_4 molten salt to conserve initial actinide mass loading but reduce the peak plutonium fraction (mol%) of the molten salt during lifetime in an effort to preclude possible plutonium solubility concerns [7]. This new study with an alternate molten salt composition was identified as case MC1 since it was the first modified composition study.

8.3.2. Methods and Model Description

Case MC1 used identical input files to case RM2, discussed above, but changed the initial molten salt composition in the MCNP model to reflect a composition of 6 mol% UF_4 and 18 mol% ThF_4 rather than the base case composition of 12 mol% UF_4 and 12 mol% ThF_4 used in case RM2. All other parameters remained unchanged.

8.3.3. Results

The results from case MC1, shown below as a series of figures, indicate acceptable performance with this alternate molten salt composition and confirm the expected reduction in peak plutonium fraction (mol%) of the molten salt.

The power curves displayed in Figure 35 show a slight reduction in operation time at plateau power for case MC1 compared to case RM2 as well as a slightly longer tail of lower power operation. These results will be quantified in detail in subsequent figures.

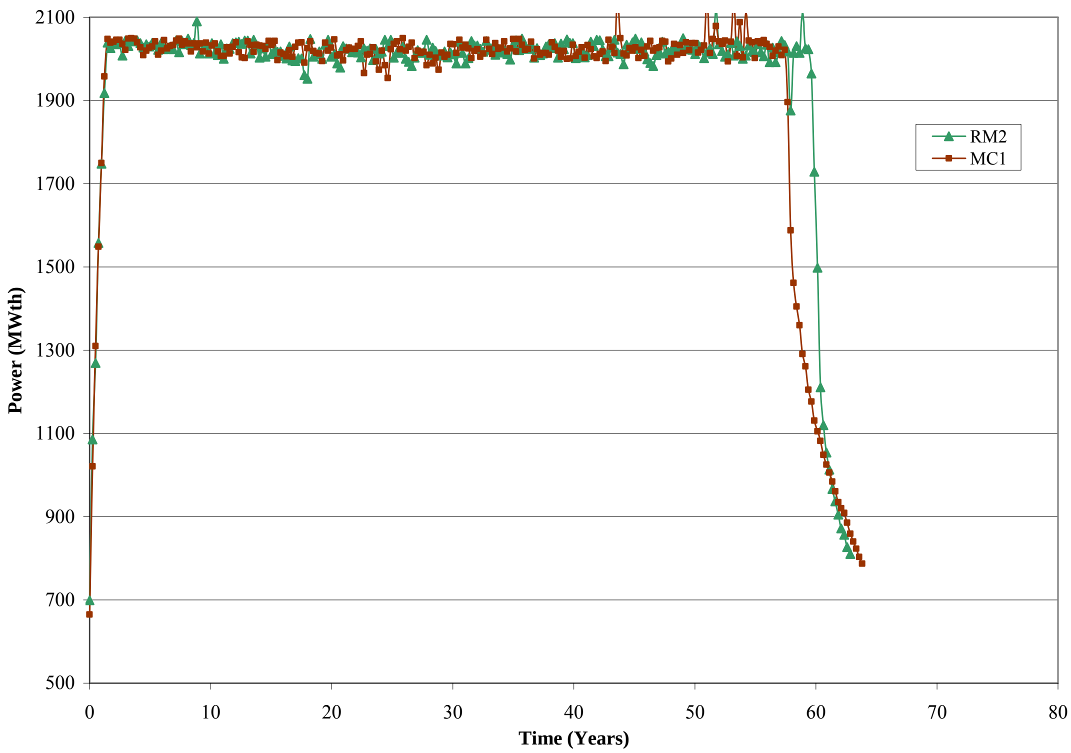

Figure 35. Power Curves for Cases RM2 and MC1

Figure 36, which shows the tritium mass curves for cases RM2 and MC1, indicates a higher peak tritium inventory for case MC1 but then a sharp drop around 56 years into operation. Based upon this figure, the power curves for these cases, and a general sense for molten salt LIFE system this drop is believed to be due to running out of fissile isotopes and having to switch to burning fissionable isotopes or more heavily convert fertile isotopes. This intuition has not been verified through any analytic means such as checking time-dependent isotopic masses of fissile isotopes but it could be at a later point if needed.

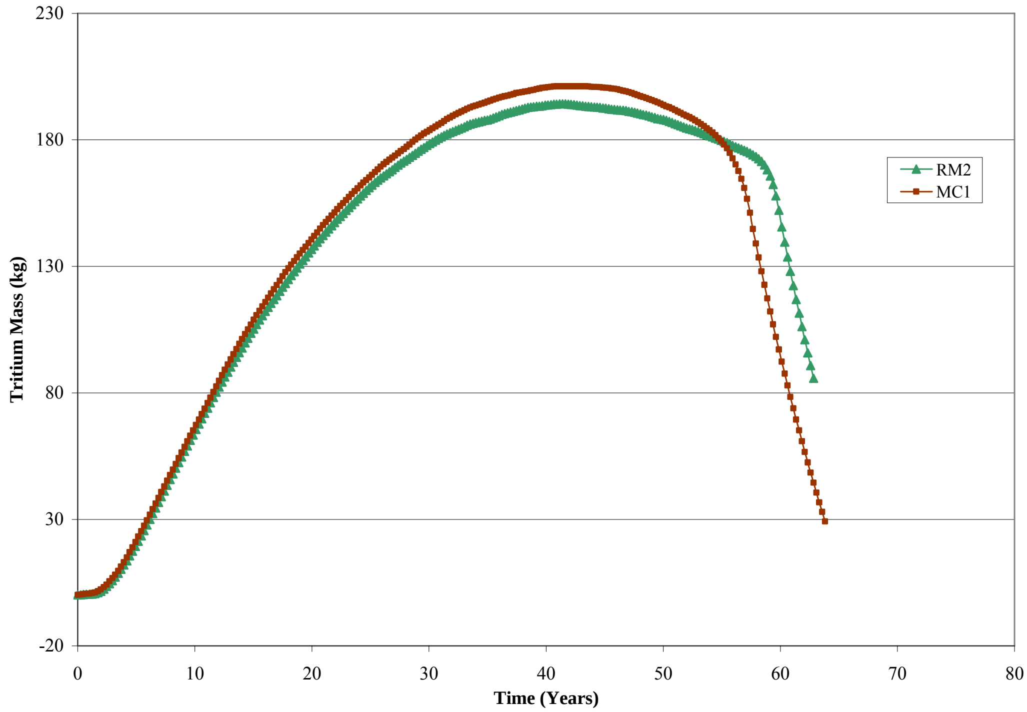

Figure 36. Tritium Mass Curves for Cases RM2 and MC1

The actinide masses curves for cases RM2 and MC1, shown below in Figure 37, indicate an incineration time of about 64 years for MC1. This is an increase in incineration time of about 1 year compared to case RM2 and is a reduction in design performance.

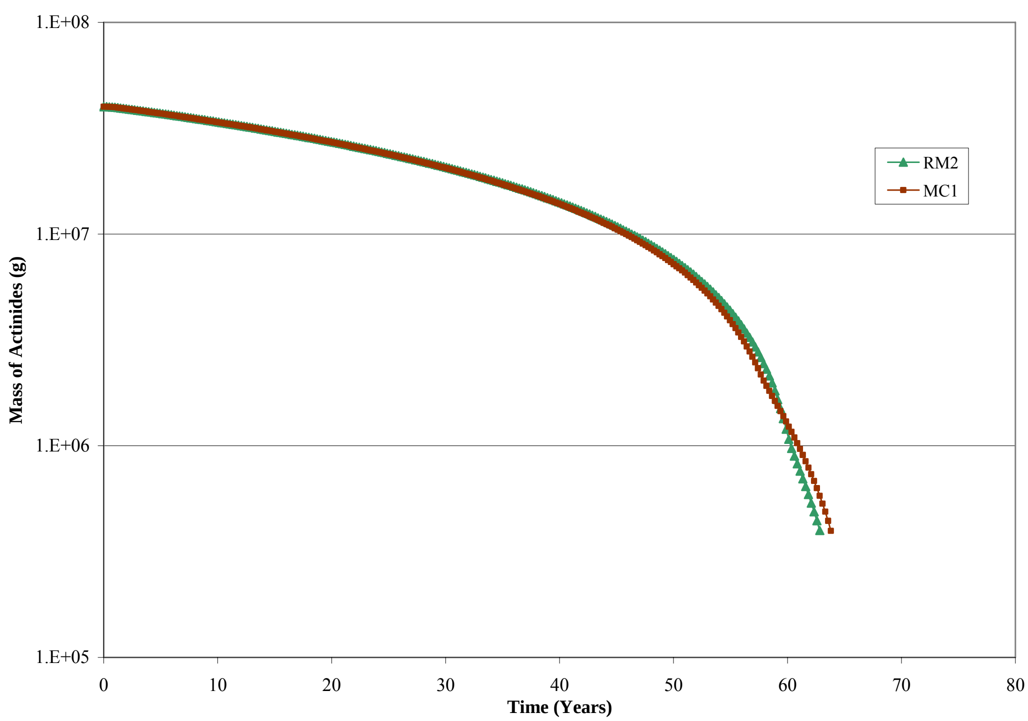

Figure 37. Actinide Mass Curves for Cases RM2 and MC1

The performance indices of MC1 indicate a reduction in design performance relative to case RM2, as seen in Figure 38.

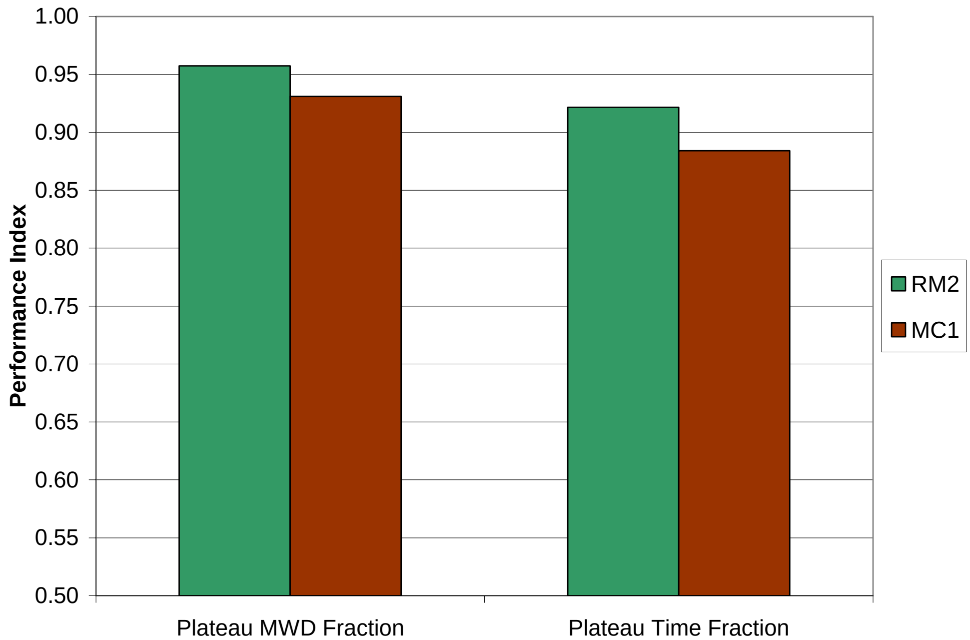

Figure 38. Performance Indices for Cases RM2 and MC1

A full analysis of time-dependent mol percent compositions and isotopic masses was not performed on the results of case MC1, but a simple calculation was made at the time of peak plutonium concentration identified in cases G70 and RM2 and yielded a peak plutonium level of about 1.2 mol%. This confirms the expected result of altering the molten salt composition and supports the goal of reducing the peak plutonium composition.

This study established the overall feasibility of using an alternate molten salt composition to attain the goals of peak plutonium levels while also achieving reasonably good system performance. There was some decrement to performance when compared to case RM2, but this reduction in performance was not of a sufficient level to preclude adequate design space existing in molten salts of very high thorium contents.

8.4. Hot Spot Ignition Design

This section documents a scoping study performed to identify available design space and initial design starting points for a molten salt blanket LIFE engine that uses a hot spot ignition (HSI) fusion system.

8.4.1. Purpose

The reference case for LIFE engine fusion driver has shifted from an ambitious fast ignition (FI) system to a NIF-like hot spot ignition (HSI) fusion drive that is believed to be more realistic for near-term development and deployment. In response to this, neutronics studies are being performed for both TRISO and molten salt blanket systems to quantify the effects of shifting to HSI parameters that include a larger fusion chamber, 1000 MW of fusion power (75MJ at a repetition rate of 13.33 Hz) instead of 500 MW, and a plateau power level of 4000 MW_{th} instead of 2000 MW_{th}. A series of studies was initiated to specifically at molten salt blanket designs for a HSI geometry and parameter set with differing amounts of initial actinide fuel loading.

Case h10 investigated a molten salt blanket with HSI driver with 40 metric tons (MT) of initial actinide loading and a plateau power of 4000 MW_{th}. Case h20 used 80 MT of initial actinide loading with a plateau power of 4000 MW_{th}, which was believed to match more directly to the molten salt fast ignition reference case G70 which had 40 MT of initial actinide loading and a plateau power level of 2000 MW_{th}. Case h32, which followed two initial attempts that mistakenly used unrepresentative input parameters, used an initial actinide loading of 120 MT and a plateau power level of 6000 MW_{th}.

Table 7 provides a clear report of the initial actinide loading in metric tons (MT), initial mass of the fuel and moderator region, initial total mass of the entire MCNP model counting all components listed in Table 2, and the plateau power level of all three hot spot ignition molten salt cases of interest.

Table 7. Masses and Plateau Power Level for Hot Spot Ignition Study Cases h10, h20, and h32

Case	Initial Actinide Loading (MT)	Initial Mass of Fuel and Moderator Region (MT)	Initial Total Mass of MCNP Model (MT)	Plateau Power Level (MWth)
h10	39.8	119.5	638.8	4000
h20	79.5	239.1	791.7	4000
h32	119.3	358.6	943.3	6000

8.4.2. Methods and Model Description

The hot spot ignition (HSI) design study used an MCNP model described in Section 6.2. The initial design was taken from case bb0, a hot spot ignition LIFE engine concept with a TRISO fuel form blanket design. Following a conversation with Ryan Abbott at LLNL regarding the structural and thermal requirements that would likely be necessary for a molten salt blanket design based upon case bb0, the geometry shown and described in Section 6.2 was arrived at.

The fuel and moderator region composition and density, along with the list of isotopes and elements undergoing continuous active removal, were taken from case RM2 described above. The fuel and moderator region thickness was set for each case to provide the desired initial actinide loading for each case using the specified molten salt density and composition.

8.4.3. Results

The results shown below indicate that sufficient design space exists to continue studying hot spot ignition LIFE engine concepts using molten salt blanket designs.

Figure 39, which shows the power curves for all three HSI cases, provides evidence that all three molten salt HSI cases perform adequately to establish feasibility. It is important to note that the variation seen in the power curve for case h32 indicates that an insufficient number of iterations in the TBR code were used (the maximum number of iterations used for case h32 was 16) and that the system may be coming closer to criticality.

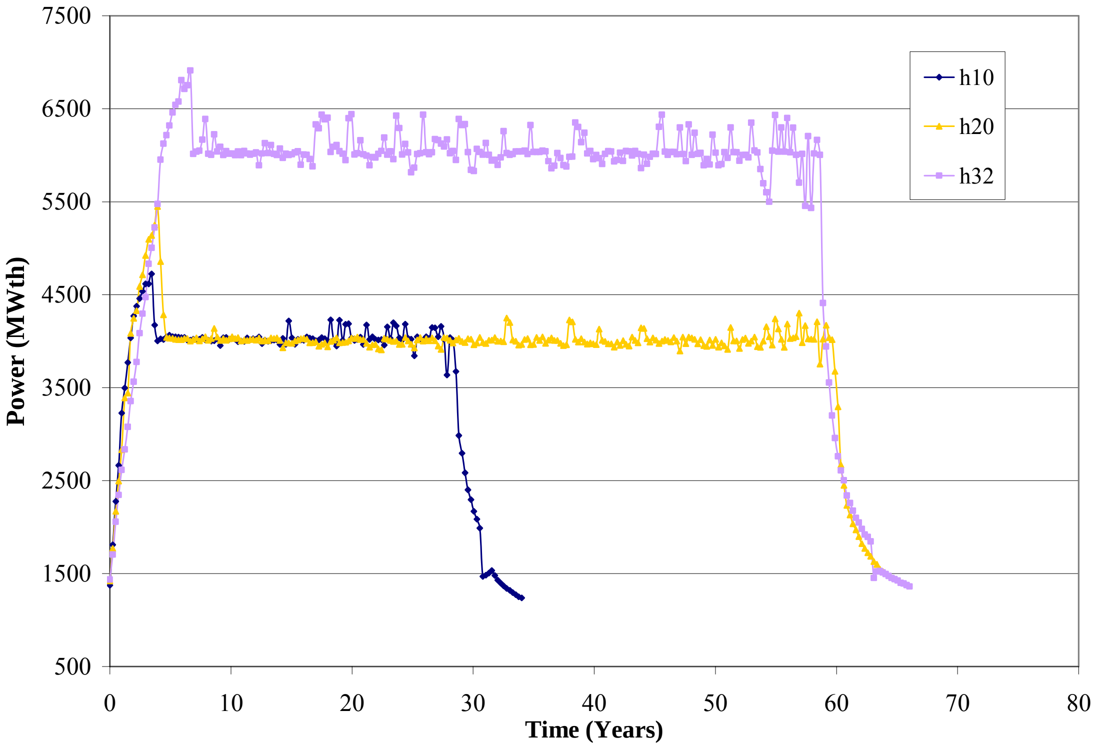

Figure 39. Power Curves for Molten Salt HSI Cases h10, h20, and h32

The TBR and tritium mass curves for the HIS cases, shown in Figure 40and Figure 41 respectively, indicate that both h10 and h32 ran out of tritium before reaching the 99% FIMA burnup goal while case h20 still had over 25kg of excess tritium inventory at 99% FIMA burnup. This implies that the plateau power level of h20 could be pushed slightly higher to yield better performance for that design. All three cases show sufficient variability in power and tritium production that they need to be re-run with tighter controls on power and TBR levels, as well as an increase in the maximum number of TBR iterations from 16 to perhaps 32, in order to accurately depict the performance possible for each case. As discussed previous for other studies, case h20 demonstrates the sharp drop in tritium mass and TBR near its end of life that indicates it was running short on fissile and fissionable material to burn before the end of the power plateau. Cases h10 and h20 also exhibit somewhat similar behavior, implying that a more careful burnup calculation where the TBR and Power ranges in the TBR code are more carefully controlled may yield options for the h10 and h32 cases that reach a better balance of reaching 99% FIMA burnup and exhausting the tritium inventory at the same time and thus improve performance.

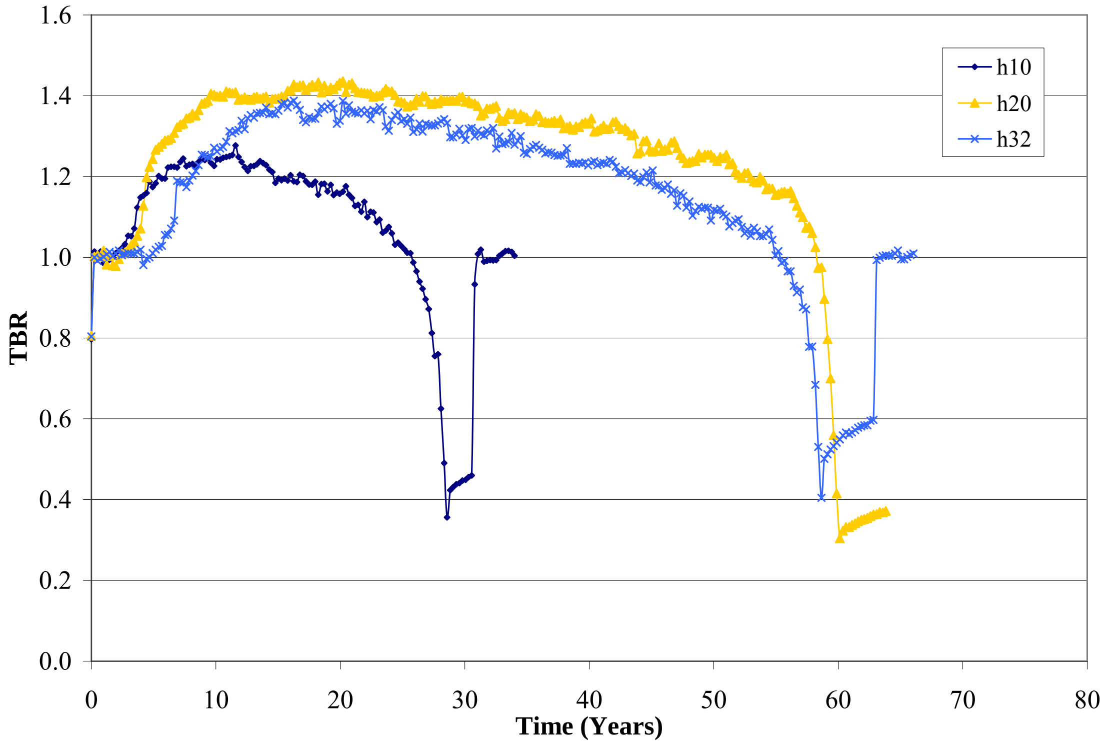

Figure 40. TBR Curves for Molten Salt HSI Cases h10, h20, and h32

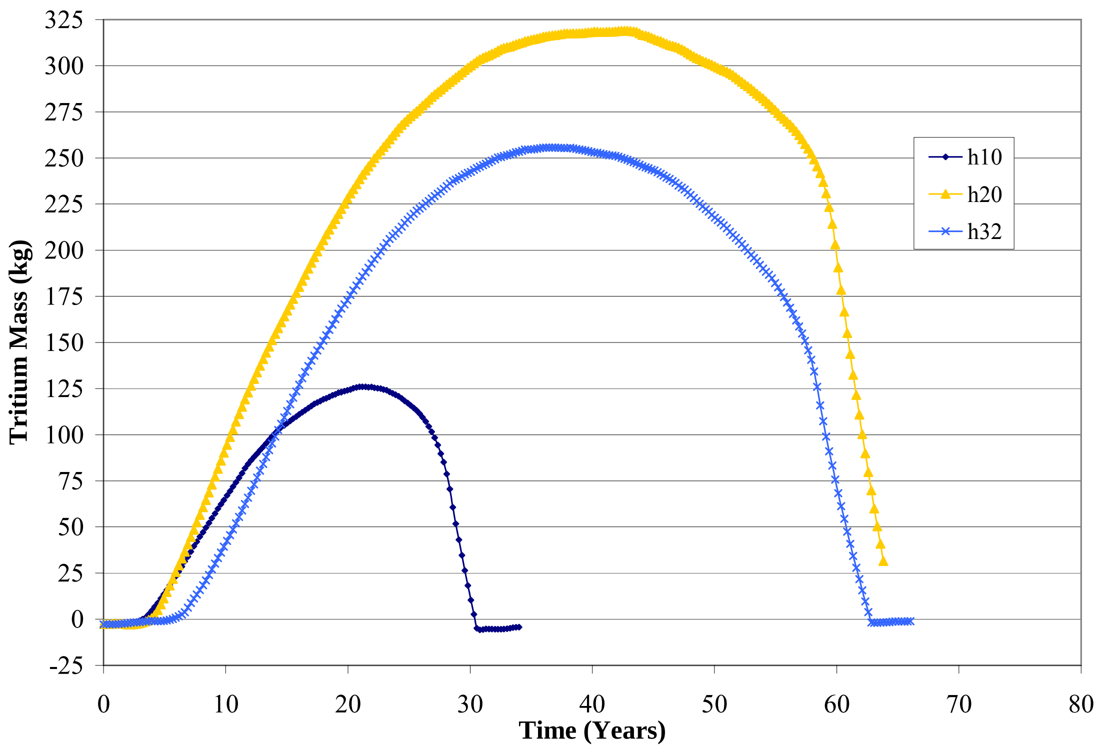

Figure 41. Tritium Mass Curves for Molten Salt HSI Cases h10, h20, and h32

The actinide mass curves in Figure 42 provide a clear picture of the different initial actinide loading levels as well as the rates at which they are consumed. The incineration time of h10 is roughly half the incineration time of each of the other cases, which was expected based upon the previously discussed actinide loading levels and plateau power levels. The slight increase in incineration time for h32 as compared to h20, a difference of about 2 years, provides some evidence that all of these cases need further optimization and that case h32 in particular is far from being fully optimized.

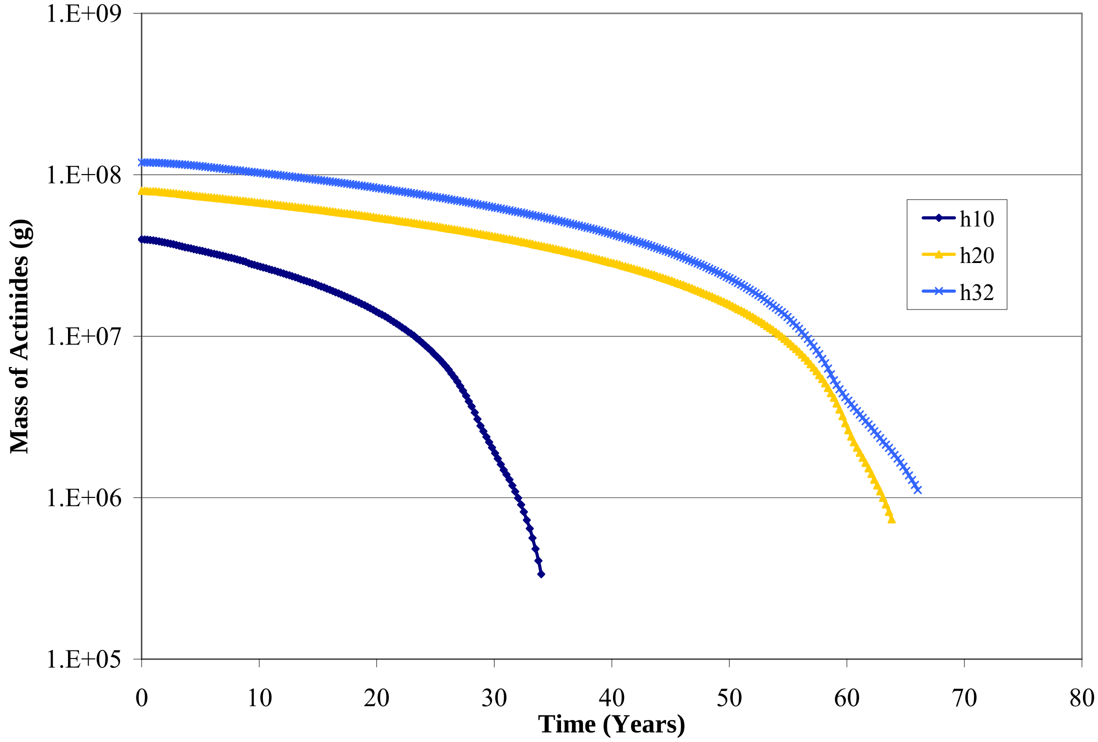

Figure 42. Actinide Mass Curves for Molten Salt HSI Cases h10, h20, and h32

The performance indices, shown below in Figure 43, provide a more concrete metric to show that the molten salt HIS cases have not been fully optimized. While some performance differences between fast ignition and hot spot ignition may be expected, the fact that case h20 has almost the same level of performance as case G70 shows that hot spot ignition systems should be able to achieve roughly the same performance as fast ignition systems if for molten salt blanket designs if the assumptions in the models for this study are realistic.

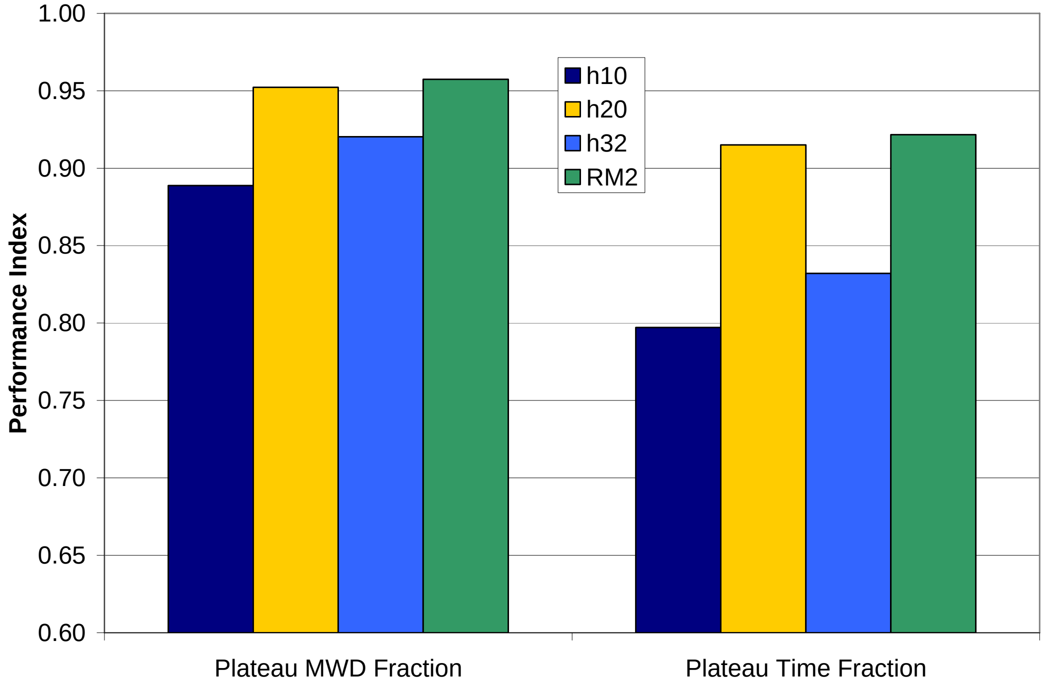

Figure 43. Performance Indices for Cases h10, h20, h32, and RM2

Based upon all of these results, compelling designs seem feasible using initial actinide loadings of both 80 metric tons and 120 metric tons. Further work should be performed to study this design space in greater detail.

9. SIDE STUDIES

This section documents several side studies that were performed during these molten salt preliminary design analyses. These side studies examine the impact of the fission product (FP) removal assumed in the main molten salt concept and investigate the design space available in reflector optimization through varying the reflector material or thickness,

9.1.Impact of No Fission Product Removal

9.1.1. Purpose

Stemming from conversations with Jeff Latkowski and others, this study was performed to investigate the effects of not removing any of the semi-noble and noble metals and gaseous fission products being continuously removed during the Molten Salt Fast Ignition Moderator Study. This scenario provides a more direct comparison to a TRISO-based concept, and provides insight on what benefits come from fission product removal.

9.1.2. Methods and Model Description

Calculations performed specifically for this study will be denoted in text and plots with a suffix "nofprm" to indicate that no fission product removal was performed by the TBR code during burnup calculations.

All “nofprm” cases used identical input files to the case they are modifying (e.g., “G80nofprm” used the input file from case G80 from the MSFIMS) other than removing the specifications for removing the noble and semi-noble metals and gaseous fission products listed in Section 4. Only tritium was marked for continuous removal in the “nofprm” series of cases. The series of “nofprm” cases were performed using cases F55, G05, G10, G20, G40, G60, and G80 from the MSFIMS. A “nofprm” case was not run for G70 because the “nofprm” calculations were performed in parallel with the MSFIMS and thus were being run before G70 was identified as the fast ignition reference case for molten salt. The general trends, however, indicate that they would hold true for G70.

9.1.3. Results

The results of this study indicate that the removal of the noble and semi-noble metals and gaseous fission products listed in Section 4 provides a significant performance benefit, as expected.

The following figures provide a comparison of two particular cases (G60 and G8) along with their “nofprm” modified versions, in an effort to present a small subset of curves that can be seen and understood clearly.

Figure 44 shows the power curves for times after 40 years of operation for this set of cases. These curves indicate that fission product removal significantly extends the period of plateau power operation.

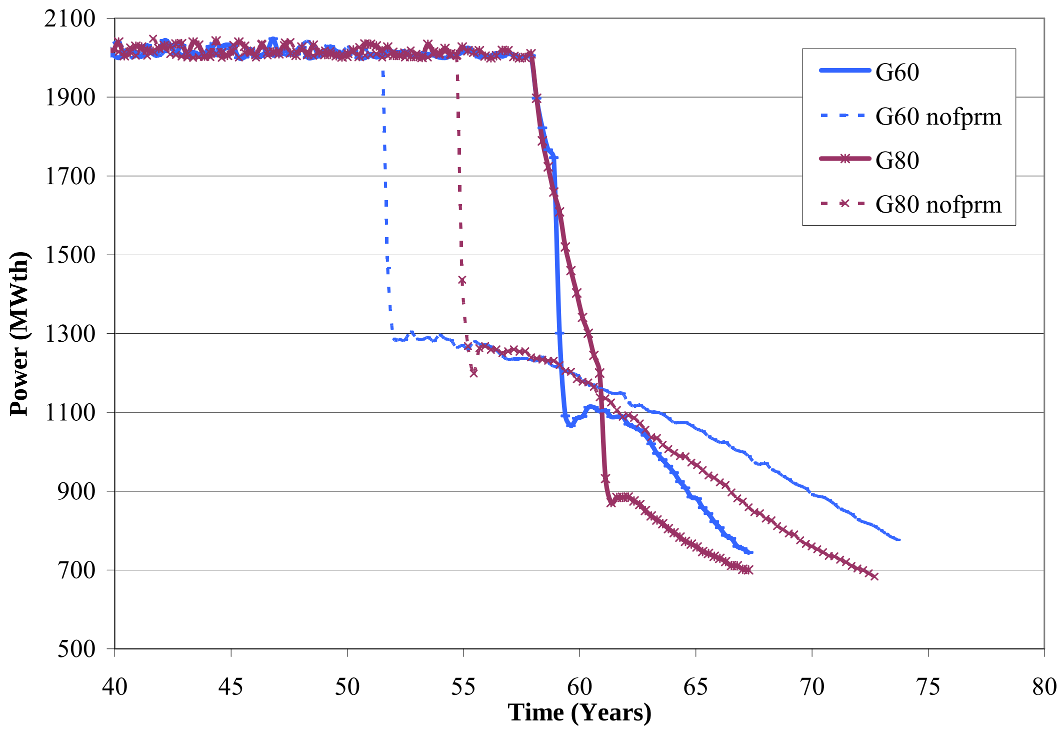

Figure 44. Power Curves for Cases G60 and G80 With and Without FP Removal

The tritium mass curves for this set of cases, found below in Figure 45, display the greatly increased tritium production possible (and needed) in cases with the fission product removal.

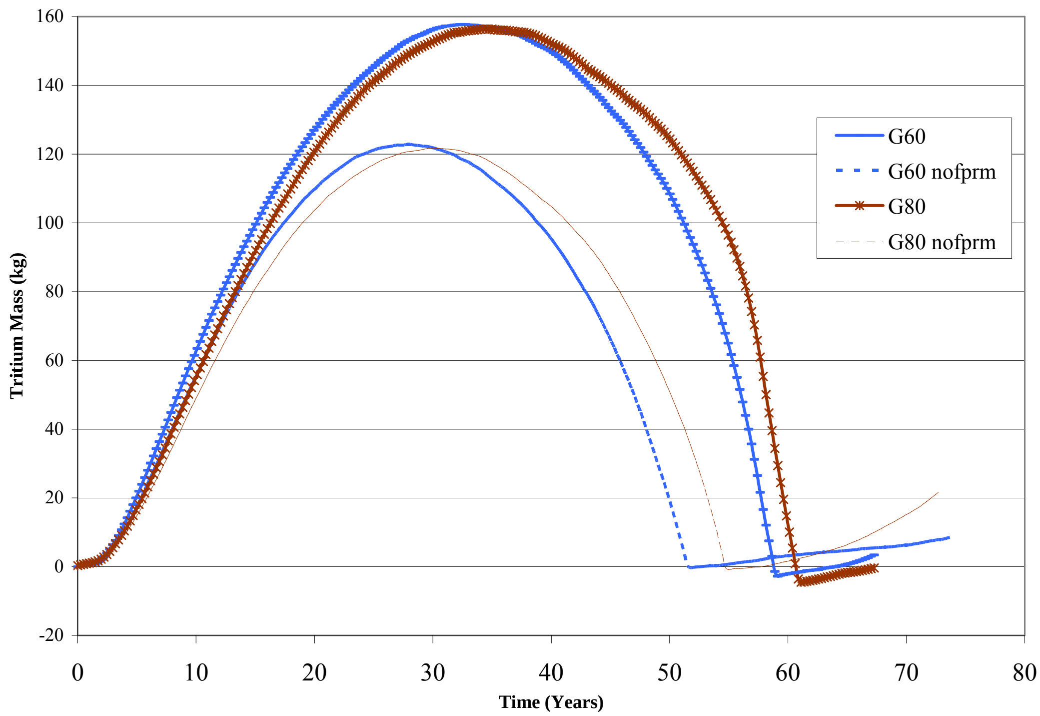

Figure 45. Tritium Mass Curves for Cass G60 and G80 With and Without FP Removal

The actinide mass curves in Figure 46 show that the required incineration time is significantly less for a case that has the noble and semi-noble metals and gaseous fission products being removed and in fact the G60 "nofprm" case did not reach 99% FIMA.

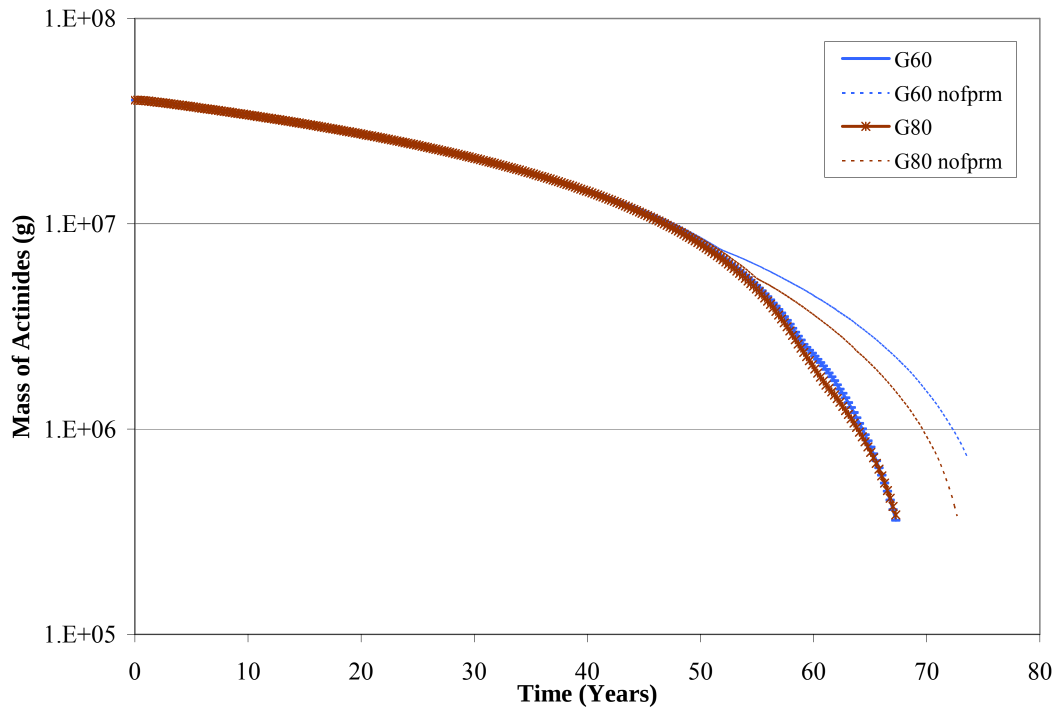

Figure 46. Actinide Mass Curves for Cass G60 and G80 With and Without FP Removal

A comparison of the performance indices from these cases can be found in Figure 47. This comparison displays the clear reduction in performance in the "nofprm" cases.

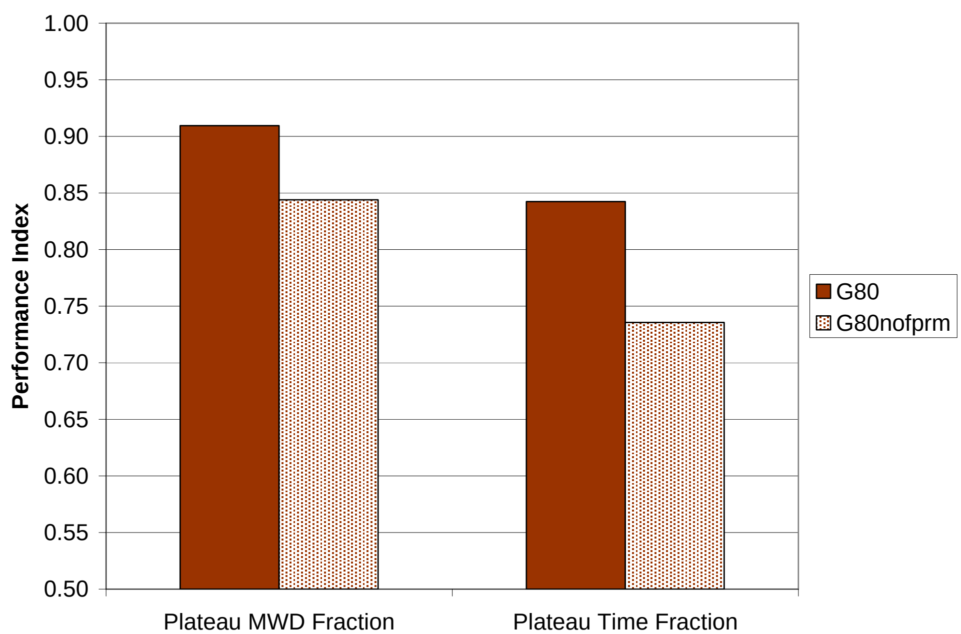

Figure 47. Comparison of Performance Indices for Case G80 With and Without FP Removal

This study has identified and quantified substantial performance improvements through removing noble and semi noble metals and gaseous fission products. The results support this continuous removal process during operation of a LIFE engine with a molten salt blanket design by showing the clear neutronics benefits. This continuous removal is likely necessary from a chemistry and materials perspective anyway, due to the potential for the metallic elements to plate out in cold parts of the system and the likelihood of the gaseous elements naturally dissociating from the molten salt.

The ebenfits identified in this study also support the exploration of other concepts

9.2. Reflector Optimization

9.2.1. Purpose

The purpose of this study was to investigate and quantify the amount of performance improvement possible through optimizing the neutron reflector. Previous studies have been performed to look at reflector optimization for blanket designs using TRISO fuel, but no studies like this have been performed in the recent past for TRISO and no such studies have been performed previously for molten salt concepts.

The first part of this study examined the design space available simply through varying the thickness of the existing neutron reflector. Graphite reflector thicknesses of 0.1, 30, 75, 100, 125, and 500 cm were examined during this portion of the reflector optimization study. This range of thicknesses was intended to capture the effects of having a reflector that ranges from essentially nonexistent to an “infinite” thickness from the perspective of neutrons in the system. In the middle were the base case reflector thickness (30cm) and thicknesses chosen to evaluate what thickness allowed most of the performance improvements of an infinite reflector while adding minimal mass and thickness to the blanket design.

The second part of the reflector optimization study investigated the impact of varying both reflector thickness and material. Beryllium oxide (BeO) at a density of 2.85 g/cc, beryllium metal (Be) at a density of 1.85 g/cc, and alumina (Al_2O_3) at a density of 3.95 g/cc were used in separate burnup calculations in place of the graphite at 2.0 g/cc used in the base case.

9.2.2. Methods and Model Description

All cases in the reflector optimization study used the input files from case G15 (molten salt with 15 vol% graphite moderator) from the molten salt fast ignition moderator study (MSFIMS) as a base case. The reflector material and thickness in the G15 input files were modified as necessary for each case. All other parameters remained unchanged except for the radii of the final wall outside the reflector, which are directly dependent on the reflector outer radius.

Case G15 was chosen as the base case because the reflector optimization study was performed in parallel with the MSFIMS and thus was being run before G70 was identified as the fast ignition reference case for molten salt. It is unknown whether these trends would hold true for G70 since G70 has significantly more moderator in the molten salt.

9.2.3. Results

The reflector optimization study determined that significant performance improvements could be realized through using a thicker graphite reflector or switching to a more efficient reflector material.

Table 8 provides a summary of the reflector study cases as well as important parameters for each in order to understand the similarities and differences of the different cases.

Table 8. Important Parameters for Reflector Study Cases

Case	Reflector Thickness (cm)	Reflector Material	Reflector Mass (MT)	BOL Portion of Total System (wt%)
G15	Graphite	30	59.03	35.1
R11	Graphite	500	3562.37	96.6
R12	Graphite	0.1	0.18	0.2
R13	Graphite	75	172.77	61.1
R14	Graphite	100	250.47	69.4
R15	Graphite	125	339.52	75.3
R16	BeO	30	84.11	43.5
R17	Al2O3	30	116.57	51.6
R18	Be	30	54.60	33.4
R19	BeO	75	483.81	81.3
R20	BeO	125	246.19	69.1
R21	Al2O3	75	670.55	85.8
R22	Al2O3	125	341.22	75.6
R23	Be	75	314.05	73.9
R24	Be	125	159.81	59.2

As stated above, the first part of the study investigated the design space available with case G15 for a wide range of graphite reflector thicknesses. The power curves from these burnup calculations are shown in Figure 48.

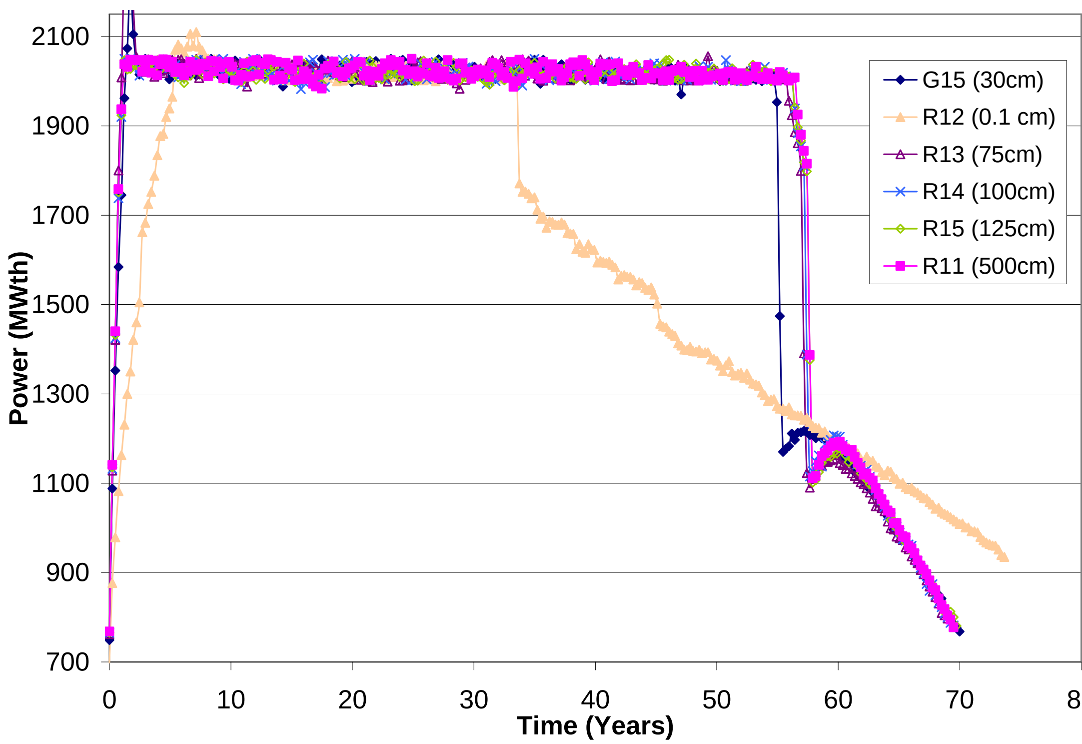

Figure 48. Power Curves for Several Graphite Reflector Thicknesses

The power curves indicate that a 30 cm graphite reflector, which is base case G15, provides a power plateau that both starts earlier in operation and lasts much longer than the results from case R12 which is essentially without a reflector. Going to thicker graphite reflectors offers a substantial benefit over a 30cm reflector in terms of extending the power plateau and ramping the power down faster during Phase 3 of operation.

The second part of the reflector optimization study examined the design space available for case G15 with other reflector material options at a range of reflector thicknesses. Figure 49 displays a subset of the power curves that results from this portion of the reflector study, demonstrating the power curves after 50 years of operation for thicknesses of 30cm and 125cm for all 4 materials considered. The results indicate that if a 30cm reflector is preferred due to size limits of the blanket design, Be and BeO reflectors offer significant improvements to the power curve over graphite due to longer operation at plateau power.

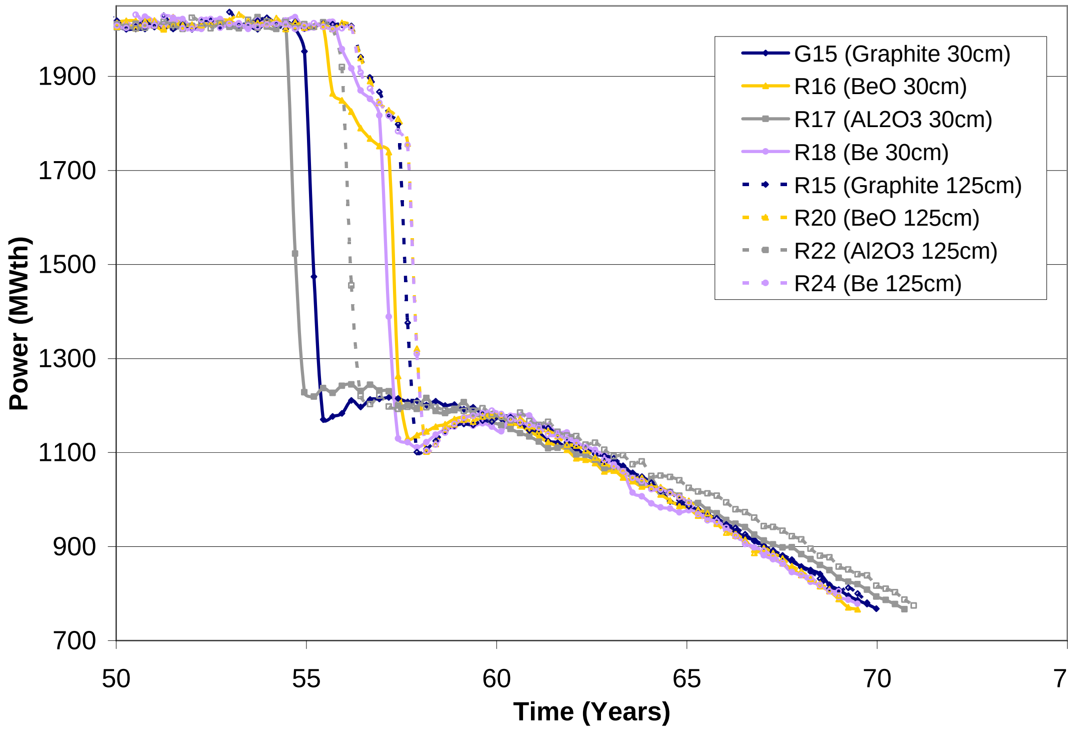

Figure 49. Power Curves for Various Reflector Materials at 30cm and 125cm Thicknesses

Figure 50 shows a comparison of the performance indices as a function of both reflector material and thickness. These results indicate that a 30cm reflector has significantly better performance with Be or BeO as the reflector material, but as the reflector thickness increases the difference between beryllium-based reflectors and graphite decreases. The beryllium-based reflector designs retain slightly higher performance indices than graphite at 75cm to 100cm but the differences become very small.

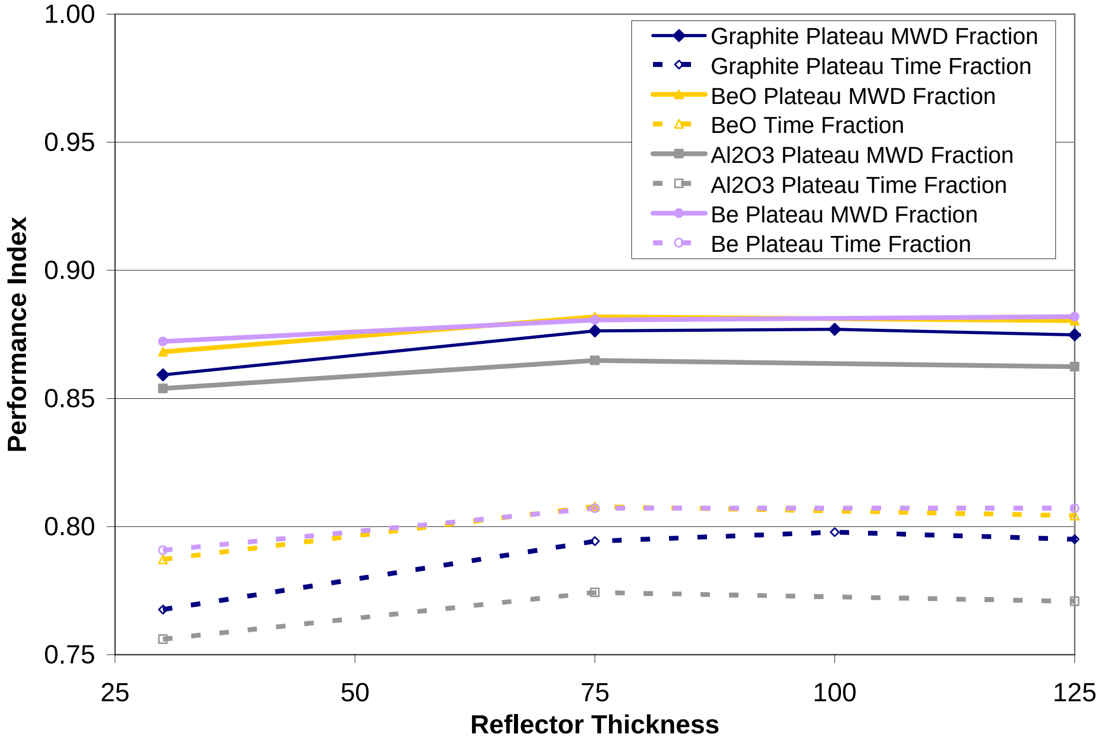

Figure 50. Comparison of Performance Indices for Varying Reflector Thicknesses and Materials

The results from this reflector optimization study indicate that if a graphite reflector is retained moving forward with other designs it should be increased to thicknesses around 75 cm. The results also indicate possible advantages of switching to Be or BeO reflectors if size is a driving limitation and possible advantages of switching to Be if mass is a driving limitation. Switching to a Be or BeO reflector would bring in new issues, however, in the form of considerations for several key areas: radiation shielding due to neutron and gamma radiation from the beryllium, environmental safety due to beryllium handling and contamination, and possible severe cost implications for using beryllium instead of graphite.

The results from this study offer possible performance improvement options to all LIFE engine blanket designs. A reflector optimization study should likely be performed specifically for future designs rather than rely entirely upon this one, however, since the neutron flux spectrum and leakage and other aspects of case G15 led to a different optimal reflector than might be true for other designs.

10. QUALITY ASSURANCE AND VERIFICATION EFFORTS

Quality assurance and verification efforts, important to all research and design efforts, have been made throughout the analyses in this report. These efforts can be distinctly broken down into how they apply to the software used and how they apply to the results obtained or derived from the software.

The two main computer codes of consequence in LIFE neutronics analyses in this report are MONTEBURNS 2.0 and TBR. The MONTEBURNS 2.0 burnup code package has been benchmarked by other users [2] and is thus deemed an acceptable code package for design work. The TBR code has been analytically verified and has undergone continual comparison to results obtained with previous versions of TBR as the code evolves in order to verify new code versions.

Initial verification efforts consisted of independently reproducing previous results obtained from studies conducted for the LIFE Program, including both burnup calculations as well as analysis of results. Cases da0 and F55 were re-run by the analyst and confirmed to match the original results previously documented within the limits of the stochastic uncertainty in MONTEBURNS. This level of verification ensured that the analyst had a proper working knowledge of the codes and analysis techniques to be used in later studies.

Verification efforts were also made to ensure that the raw results coming from the MONTEBURNS and TBR calculations, as well as all derived values and parameters arrived at through analysis of raw results, are reasonable. Actions taken as part of these verification efforts include spot checking of time-dependent values of isotopic masses and concentrations, overall masses, power curves and TBR curves, and searching through output files to look for any error messages from the computer codes themselves. Output files from the MONTEBURNS and TBR process specifically include lines to indicate the amount of mass in the problem "lost" due to truncation errors or missing cross sections. All results reported in this report passed this basic quality assurance process unless otherwise noted.

11. CONCLUSIONS

Preliminary studies into LIFE engines with molten salt fueled blanket designs prove the feasibility of concepts with fusion driver systems based upon both fast ignition and hot spot ignition. The optimal composition of the molten salt and graphite region was found to be 70 vol% graphite and 30 vol% molten salt.

The recommended molten salt fast ignition reference case from the moderator study was case G70, though subsequent analyses show that RM2 would provide even better performance and should be used for further studies. This case has an initial actinide loading of 40 metric tons with equal amounts of depleted uranium and thorium. The plateau power level is 2000 MW_{th} and 63 years of operation were required to reach a burnup of 99% FIMA. Alternate molten salt compositions may be needed to keep peak plutonium levels below solubility limits.

The recommended molten salt hot spot ignition reference case is h20, which has an initial actinide loading of 80 metric tons with equal amounts of depleted uranium and thorium. The plateau power level is 4000 MW_{th} and 64 years of operation were required to reach a burnup of 99% FIMA. Alternate molten salt compositions may again be needed to keep peak plutonium levels below solubility limits.

Studies performed to examine the impact of various levels of fission product removal show that design performance becomes better as more fission products are removed.

Reflector optimization studies showed that a graphite reflector of about 75cm to 100cm could provide significant performance improvements over the existing 30cm graphite reflector.

12. FUTURE WORK

During the course of performing the various studies and analyses in this report, as well as interpreting those studies and drawing conclusions from them, numerous areas of further study have been identified for future work on LIFE engine blanket designs of all types and molten salt specifically. General items applicable to all LIFE neutronics work include optimization studies of the reflector and beryllium multiplier regions, evaluating the effects of using explicit heterogeneous geometries in the MCNP models, and adding reaction rate calculations as an output of the TBR code. Molten salt specific items include radiation damage calculations for graphite and structural materials, further development of hot spot ignition concepts, and further investigation of concepts that remove fission products and perhaps even have continuous feed of new fuel.

The reflector optimization study in Section 9.2 points to the gains that can be made through optimizing not only the reflector but other components of the system. Future work should continue to investigate reflector optimization through investigating both different thicknesses and different material options, especially as the relate to cases such as G70 and RM2 that have a much higher moderator volume percentage in the molten salt region than case G15 used as a base case for the reflector study in this report. Optimization studies should also be expanded to look at optimization of the beryllium neutron multiplier region

As was mentioned in Section 6, all work to date has been performed assuming homogenous regions in the MCNP model. Future work should look at the effect of moving from homogenous modeling to various levels of explicit heterogeneous geometry models. Existing work is already being performed in this area to investigate different levels of heterogeneity in TRISO fuel cases, and scoping level work for single time steps has been performed to investigate the mechanics of how to include the graphite in the molten salt region as distinct plates or shells of graphite with molten salt between them, but this work should be expanded for both TRISO and molten salt options to see the lifetime effect of explicit geometry.

Reaction rate calculations for various reactions of interest (e.g., neutron capture and fission reaction rates for U-238 and Pu-239 as well as tritium production from Li-6 and Li-7) should be added to the existing process so that they can be visualized during post-processing and enable a more fundamental understanding of the system. During a brief side study not documented in this report, it was demonstrated that the TBR and neutron capture and fission reaction rates could be directly calculated using a combination of modifications to the MCNP input and modification to the TBR code. The MCNP input must be modified to output the desired flux-weighted microscopic cross sections through the use of a Tally Multiplier card, and then these flux-weighted cross-sections can be combined with parameters already calculated in the TBR code to provide reaction reactions and the TBR of the MCNP model at the level of a single cell or for the entire system. These reaction rates and TBR calculations could also be independently reported for each material if desired. The TBR calculation performed during this study reproduced the TBR reported from the TBR code to within 0.5%, which is a deviation that could be due to stochastic uncertainties or differences in nuclear cross section data used or a combination of both factors.

Radiation damage calculations are needed for all LIFE blanket designs, but molten salt designs negate the concerns about radiation damage of fuel and instead narrow the focus onto the radiation damage predicted in graphite moderator and reflector regions as well as the structural materials. These issues, which involve both neutronics and materials science and engineering, need to be resolved to determine the expected lifetime of graphite and structural components. If the lifetime of graphite becomes an issue due to structural concerns, analyses could be performed to investigate the possibility of using a steel or other structural material to can the graphite and provide structural support for it. Specific analysis should also be performed to determine the level of corrosion in the molten salt system and how this impacts other design parameters.

Hot spot ignition, which has been identified as the baseline design LIFE engines, was studied briefly for molten salt as documented in Section 8.4 but should be investigated further to find optimized designs and establish bounds on the total design space. The results in Section 8.4 establish some elements that show the basic feasibility of

Lastly, options should be investigated to look at other molten salt compositions (as shown in Section 8.3) as well as different schemes of fission product removal (as documented in Section 8.2). Ralph Moir and others have also raised the possibility of moving to concepts that involve continuous feed of some level of fresh fuel to the molten salt in order to maintain a constant plateau power level indefinitely or constant removal of fissile material in order to use a molten salt LIFE blanket as a breeding platform for nuclear fission reactors. Even within studies of the same molten salt composition and without continuous feed or fresh fuel, Section 8.2 suggested further studies which increase the plateau power level of case RM3 even further. All of these options merit further study when considering molten salt LIFE engine blankets.

13. REFERENCES

1. J.F. Latkowski et al. “LIFE: Laser Inertial Fusion/Fission Energy,” awaiting publication.

2. D.L. Poston, H.R. Trellue. "User's Manual, Version 2.0 for MONTEBURNS Version 1.0," LA-UR-99-4999, (September 1999).

3. X-5 Monte Carlo Team, "MCNP - A General Monte Carlo Code for N-Particle Neutron Transport, Version 5 - Volume I, Overview and Theory," LA-UR-03-1987, Los Alamos National Laboratory. April 2003.

4. J.J. Powers. "Summary of Neutronics Depletion (‘Burnup’) Code Survey Assessment," internal LLNL LIFE Neutronics memo, (August 6, 2008).

5. K.J. Kramer. "Algorithm for computing Tritium Breeding Ratio," internal LLNL LIFE Neutronics memo, (August 2007).

6. R.A. Schwarz, MCNP Visual Editor, http://mcnpvised.com.

7. K. Furukawa, H. Ohno. "Molten LiF-BeF_2 (Flibe) System," Volume 1 of Data Books for Molten Materials, Japan Nuclear-Energy Information Center, (1980).